

PATENT COOPERATION TREATY

PCT

NOTIFICATION OF ELECTION

(PCT Rule 61.2)

From the INTERNATIONAL BUREAU

To:

Assistant Commissioner for Patents
United States Patent and Trademark
Office
Box PCT
Washington, D.C.20231
ETATS-UNIS D'AMERIQUE

in its capacity as elected Office

Date of mailing (day/month/year)
23 March 2000 (23.03.00)

International application No.
PCT/GB99/02449

Applicant's or agent's file reference
A25664 WO

International filing date (day/month/year)
27 July 1999 (27.07.99)

Priority date (day/month/year)
05 August 1998 (05.08.98)

Applicant

OATES, Martin, John

1. The designated Office is hereby notified of its election made:

☒ in the demand filed with the International Preliminary Examining Authority on:
24 February 2000 (24.02.00)

☐ in a notice effecting later election filed with the International Bureau on:

2. The election ☒ was

☐ was not

made before the expiration of 19 months from the priority date or, where Rule 32 applies, within the time limit under Rule 32.2(b).

The International Bureau of WIPO
34, chemin des Colombettes
1211 Geneva 20, Switzerland

Facsimile No.: (41-22) 740.14.35

Authorized officer

Pascal Piriou

Telephone No.: (41-22) 338.83.38

09/743520

PATENT COOPERATION TREATY

RECD 09 NOV 2000

PCT

PCT

5640

INTERNATIONAL PRELIMINARY EXAMINATION REPORT

(PCT Article 36 and Rule 70)

Applicant's or agent's file reference A25664 WO	FOR FURTHER ACTION See Notification of Transmittal of International Preliminary Examination Report (Form PCT/IPEA/416)	
International application No. PCT/GB99/02449	International filing date (day/month/year) 27/07/1999	Priority date (day/month/year) 05/08/1998
International Patent Classification (IPC) or national classification and IPC G06F17/30		
Applicant BRITISH TELECOMMUNICATIONS PLC et al.		



1. This international preliminary examination report has been prepared by this International Preliminary Examining Authority and is transmitted to the applicant according to Article 36.
2. This REPORT consists of a total of 5 sheets, including this cover sheet.

☒ This report is also accompanied by ANNEXES, i.e. sheets of the description, claims and/or drawings which have been amended and are the basis for this report and/or sheets containing rectifications made before this Authority (see Rule 70.16 and Section 607 of the Administrative Instructions under the PCT).

These annexes consist of a total of 15 sheets.

3. This report contains indications relating to the following items:

 - I ☒ Basis of the report
 - II ☐ Priority
 - III ☐ Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
 - IV ☐ Lack of unity of invention
 - V ☒ Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
 - VI ☐ Certain documents cited
 - VII ☒ Certain defects in the international application
 - VIII ☐ Certain observations on the international application

Date of submission of the demand 24/02/2000	Date of completion of this report 06.11.2000
Name and mailing address of the international preliminary examining authority:  European Patent Office D-80298 Munich Tel. +49 89 2399 - 0 Tx: 523656 epmu d Fax: +49 89 2399 - 4465	Authorized officer Oestergaard, M Telephone No. +49 89 2399 2551 

**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT**

International application No. PCT/GB99/02449

I. Basis of the report

1. This report has been drawn on the basis of (*substitute sheets which have been furnished to the receiving Office in response to an invitation under Article 14 are referred to in this report as "originally filed" and are not annexed to the report since they do not contain amendments.*):

Description, pages:

2-58 as originally filed

1,59 as received on 16/10/2000 with letter of 12/10/2000

Claims, No.:

1-41 as received on 16/10/2000 with letter of 12/10/2000

Drawings, sheets:

1-30 as originally filed

2. The amendments have resulted in the cancellation of:

- ☐ the description, pages:
☐ the claims, Nos.:
☐ the drawings, sheets:

3. ☐ This report has been established as if (some of) the amendments had not been made, since they have been considered to go beyond the disclosure as filed (Rule 70.2(c)):

4. Additional observations, if necessary:

**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT**

International application No. PCT/GB99/02449

V. Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	1-41
	No:	Claims	
Inventive step (IS)	Yes:	Claims	1-41
	No:	Claims	
Industrial applicability (IA)	Yes:	Claims	1-41
	No:	Claims	

2. Citations and explanations

see separate sheet

VII. Certain defects in the international application

The following defects in the form or contents of the international application have been noted:

see separate sheet

Section V:

- 1 The claims include independent claim 1 defining a data processing apparatus and independent claim 21 defining a processor implemented method with steps corresponding to the means of claim 1. Claims 40 and 41 define sets of claims corresponding to the method claims 21 to 39. The claims can be derived from the original set of claims.

The nearest prior art document was found to be D1 = GOLDBERG D E: "Genetic Algorithms in Search, Optimization, and Machine Learning", January 1989, ADDISON-WESLEY USA. This document discloses various algorithms relating to the field of genetics. The algorithms are such as reproduction, mutation and crossover between selected individuals to produce child strings of values, see D1, page 71. For crossover operators, see D1, page 119. Using a cross point all genes on a respective side of the cross point swap from one side of the cross point to another.

As a difference to the teachings of D1 the subject matter of claim describes a scheme whereby a number of genes within an initial string of genes are selected, the number of genes and start position of the selection being selected randomly. This selection is used to replace a sequence of values of the same length as the selection in a copy of the initial string of genes, where the start position of the replacement is selected randomly. Therefore in the subject matter of claim 1, both the number of genes and the selection of a position from which the genes are selected is random rather than the selection of a particular cross point as in document D1. In claim 1 also the position into which the selected sequence is placed is selected at random as opposed to the cross point of document D1.

In view of the above and the available prior art, it appears that the current set of claims is novel, involves an inventive step and is industrially applicable in accordance with the requirements of Article 33 PCT.

Section VII:

- 2 The technical features mentioned in the claims should preferably have been followed by reference signs relating to such features, see Rule 6.2(b) PCT.

**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT - SEPARATE SHEET**

International application No. PCT/GB99/02449

- 2.1 The statement on page 59 of the description referring to a "spirit" of the claims is obviously irrelevant and should have been removed from the description in accordance with Rule 9.1(iv) PCT.

JC07 Rec'd PCT/PTO 11 JAN 2001

1

DATA PROCESSING APPARATUS AND METHOD FOR
OPTIMISING CONFIGURATION PARAMETERS
OF A PHYSICAL SYSTEM

5 The present invention generally relates to the data processing method and apparatus for determining optimum parameters of the model of a physical system and for controlling the configuration of the physical system.

10 The generation of a model of a physical system is used for a wide variety of applications including learning more about the behaviour of the system and controlling the parameters of the system. The model of the physical system can be used to try out parameters in order to determine optimum parameters which can then be
15 used for controlling the physical system. In this way the effect of the parameters on the physical system can be predicted so that only optimum parameters are chosen for use in the control of the physical system.

20 One such physical system for which configuration parameters are required is a distributed database. Because the use of corporate intranets continues to rise, the management of efficient access to data is becoming a key issue. Large information systems, often accessed on a global basis, are increasingly being provided by
25 distributed means with use of "mirror sites" to ease congestion and improve local access. The configuration of the information system defining the location of the

load on the servers and loads on the links interconnecting the servers and clients.

Although the present invention has been described hereinabove with reference to allelic representations of the solution vector, the present invention is not limited
5 to such representations and is applicable to canonical representations.

The present invention can be implemented as a computer program operating on a standard computer and
10 thus the present invention can be embodied as a storage medium containing processor implementable instructions for controlling a processor to carry out the method. Further, since the computer code can be transmitted in electronic form for example by being downloaded over a
15 network, the present invention can be embodied as an electronic signal carrying computer code for controlling a computer to carry out the method.

Although the present invention has been described hereinabove with reference to specific embodiments, the
20 present invention is not limited to such embodiments and it would be apparent to a skilled person in the art that modifications can be made within the spirit and scope of the appended claims.

CLAIMS:

1. Data processing apparatus for determining optimum parameters of a model of a physical system, the apparatus
5 comprising:

obtaining means for obtaining at least one initial string of values representing the parameters of the model to be optimised;

10 first determining means for determining a cost value associated with the model having parameters represented by the or each string of values;

generating means for repeatedly generating a new string of values by selecting a sequence of values of random length, starting at a random position in a said
15 string of values, replacing a sequence of values of the same length in a said string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values;

20 wherein said determining means is adapted to determine a cost value associated with the model having parameters represented by the new string of values; and

second determining means for determining the optimum parameters as one of said initial or new string of values
25 for which the cost value associated with the model having the optimum parameters is closest to a target.

2. Data processing apparatus according to claim 1 wherein said generating means is adapted to consider a last value in a said string of values as being sequential to a first value in said string of values such that said
5 selected sequence of values can include said last and first values sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

10

3. Data processing apparatus according to claim 1 or claim 2, wherein said obtaining means is adapted to randomly obtain values for said at least one initial string of values.

15

4. Data processing apparatus according to any preceding claim wherein said obtaining means is adapted to obtain a plurality of said initial strings of values to form a population, and said generating means is adapted to
20 perform a genetic algorithm on said population by repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in a first said string of values,

25

(b) replacing a sequence of values of the same length in a second said string of values at a random position, and

(c) mutating the value of one or more of the values of the resulting string of values to generate said new string of values.

5 5. Data processing apparatus according to claim 4, wherein said generating means is adapted to repeatedly

 (a) delete a proportion of the population for which the cost values are furthest from said target, and

 (b) perform said repeated generation of new strings
10 of values until the population is regenerated; and

 to select said first and second strings of values randomly from the undeleted proportion of the population for which the cost values are closest to said target.

15

6. Data processing apparatus according to claim 4, wherein said generating means is adapted to repeatedly

 (a) randomly select three said strings of values from said population,

20 (b) select two of the selected strings of values for which the cost values are closest to the target as said first and second strings of values, and

 (c) replace the third of said selected strings of values with the generated new string of values.

25

7. Data processing apparatus according to any one of claims 1 to 3, wherein said obtaining means is adapted

to obtain a single said initial string of values as a parent string of values and said generating means is adapted to repeatedly generating a new string of values by

5 (a) selecting a sequence of values of random length starting at a random position in said parent string of values,

 (b) replacing a sequence of values of the same length in said parent string of values at a random
10 position, and

 (c) changing the value of one or more of the values of the resulting string of values to generate said new string of values.

15 8. Data processing apparatus according to claim 7, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values if the cost value associated with the model having said new string of parameters is closer to said
20 target.

 9. Data processing apparatus according to claim 8 wherein said generating means is adapted to repeatedly replace said parent string of values with said new string
25 of values if the exponential of the difference in the cost values for said parent string of values and said new string of values divided by a factor dependent upon the

number of repetitions by said generating means is greater than a random number between 0 and 1.

10. Apparatus for controlling the configuration of a physical system, the apparatus comprising:

obtaining means for obtaining at least one initial string of values representing configuration parameters of the physical system;

modelling means for modelling the physical system using said configuration parameters to determine a cost value for the or each initial string of values;

generating means for repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values; wherein said modelling means is adapted to determine a cost value for the configuration parameters represented by said new string of values;

determining means for determining optimum configuration parameters represented as one of said initial or new string of values for which the cost value is closest to a target; and

control means for controlling the configuration of said physical system in accordance with the determined optimum configuration parameters.

5 11. Apparatus according to claim 10, wherein said generating means is adapted to consider a last value in a said string of values as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values
10 sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

12. Apparatus according to claim 10 or claim 11, wherein
15 said obtaining means is adapted to randomly obtain values for said at least one initial string of values.

13. Apparatus according to any one of claims 10 to 12, wherein said obtaining means is adapted to obtain a
20 plurality of said initial strings of values to form a population, and said generating means is adapted to perform a genetic algorithm on said population by repeatedly generating a new string of values by

(a) selecting a sequence of values of random length
25 starting at a random position in a first said string of values,

(b) replacing a sequence of values of the same length in a second said string of values at a random position, and

(c) mutating the value of one or more of the values
5 of the resulting string of values to generate said new string of values.

14. Apparatus according to claim 13, wherein said generating means is adapted to repeatedly

10 (a) delete a proportion of the population for which the cost values are furthest from said target, and

(b) perform said repeated generation of new strings of values until the population is regenerated; and

15 to select said first and second strings of values randomly from the undeleted proportion of the population for which the cost values are closest to said target.

15. Apparatus according to claim 13, wherein said
20 generating means is adapted to repeatedly

(a) randomly select three said strings of values from said population,

(b) select two of the selected strings of values for which the cost values are closest to the target as
25 said first and second strings of values and

(c) replace the third of said selected strings of values with the generated new string of values.

16. Apparatus according to any one of claims 10 to 12, wherein said obtaining means is adapted to obtain a single said initial string of values as a parent string of values and said generating means is adapted to repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in said parent string of values,

(b) replacing a sequence of values of the same length in said parent string of values at a random position, and

(c) changing the value of one or more of the values of the resulting string of values to generate said new string of values.

15

17. Apparatus according to claim 16, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values if the cost value associated with the model having said new string of parameters is closer to said target.

20

18. Apparatus according to claim 16, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values if the exponential of the difference in the cost values for said parent string of values and said new string of values divided by a factor dependent upon the number of

25

repetitions by said generating means is greater than a random number between 0 and 1.

19. Apparatus according to any one of claims 10 to 18 including means for receiving operating parameters from said physical system; wherein said model means is adapted to model said physical system using the received operating parameters.

10 20. A distributed database controller for configuring a distributed database comprising a plurality of data servers connected over a network, each data server holding any number of data units, and a plurality of clients connected over the network to the data servers, 15 each of said clients being adapted to retrieve data units from and/or update data units at one or more of the data servers, the controller comprising:

means for obtaining at least one initial string of values representing configuration parameters for the 20 distributed database defining which data server is to be accessed by which client;

modelling means for modelling the distributed database using the configuration parameters and information on the passage of data within the distributed 25 database to determine a performance value for the or each initial string of values;

generating means for repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position and changing the value of one or more of the values of the resulting string of values to generate a new string of values, wherein said modelling means is adapted to determine a performance value for the configuration parameters represented by said new string of values;

determining means for determining optimum configuration parameters represented as one of said initial or new string of values for which the performance value is the best; and

control means for configuring the distributed database to control which data server is to be accessed by which client in accordance with the determined optimum configuration parameters.

20

21. A controller according to claim 20, wherein said generating means is adapted to consider a last value in a said string of values as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values sequentially in a said string of values and the sequence

25

of values to be replaced can include said last and first values sequentially in a said string of values.

22. A controller according to claim 20 or claim 21,
5 wherein said obtaining means is adapted to randomly obtain values for said at least one initial string of values.

23. A controller according to any one of claims 20 to
10 22, wherein said obtaining means is adapted to obtain a plurality of said initial strings of values to form a population, and said generating means is adapted to perform a genetic algorithm on said population by repeatedly generating a new string of values by

15 (a) selecting a sequence of values of random length starting at a random position in a first said string of values,

(b) replacing a sequence of values of the same length in a second said string of values at a random
20 position, and

(c) mutating the value of one or more of the values of the resulting string of values to generate said new string of values.

25 24. A controller according to claim 23, wherein said generating means is adapted to repeatedly

(a) delete a proportion of the population for which the performance values are worst, and

(b) perform said repeated generation of new strings of values until the population is regenerated; and

5 to select said first and second strings of values randomly from the undeleted proportion of the population for which the performance values are best.

25. A controller according to claim 23, wherein said
10 generating means is adapted to repeatedly

(a) randomly select three said strings of values from said population,

(b) select two of the selected strings of values for which the performance values are worst, and

15 (c) replace the third of said selected strings of values with the generated new string of values.

26. A controller according to any one of claims 20 to 22, wherein said obtaining means is adapted to obtain a
20 single said initial string of values as a parent string of values and said generating means is adapted to repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in said parent string of
25 values,

(b) replacing a sequence of values of the same length in said parent string of values at a random position, and

(c) changing the value of one or more of the values
5 of the resulting string of values to generate said new string of values.

27. A controller according to claim 26, wherein said
10 generating means is adapted to repeatedly replace said parent string of values with said new string of values if the performance value associated with the model having said new string of parameters is best.

28. A controller according to claim 26, wherein said
15 generating means is adapted to repeatedly replace said parent string of values with said new string of values if the exponential of the difference in the performance values for said parent string of values and said new string of values divided by a factor dependent upon the
20 number of repetitions by said generating means is greater than a random number between 0 and 1.

29. A controller according to any one of claims 20 to
25 28, wherein said control means is adapted to copy, move and/or delete data units at data servers to change the distribution of the data units across the data servers and/or to change the data servers to be accessed by said

clients using the determined optimum configuration parameters to improve the system performance.

5 30. A controller according to any one of claims 20 to 29; wherein said modelling means is adapted to determine said performance value in dependence upon the time taken for one or more said clients to retrieve and/or update data units at one or more data servers.

10 31. A controller according to claim 30 wherein said modelling means is adapted to determine said performance value as the longest response time for any said client to access any said data server.

15 32. A controller according to claim 30, wherein said modelling means is adapted to determine said performance value as the longest response time for any said client to access any said data server and a proportion of the average of the response times for all said data servers.

20 33. A controller according to claim 30, wherein said modelling means is adapted to determine said performance value as a function of the rates at which data can be retrieved by said clients from said data servers, the
25 rates at which data can be updated by said clients at said data servers, the contention between said clients

for accessing said data servers, and communications times between said clients and said data servers.

34. A controller according to any one of claims 20 to
5 33 including monitoring means for monitoring the passage of data over the network, wherein said modelling means is adapted to use the output of said monitoring means to determine said performance values, and said control means is adapted to adaptively configure the distributed
10 database.

35. A distributed processing controller for configuring a distributed processing system comprising a plurality of servers connected over a network, each server being
15 capable of carrying out any number of processing operations, and a plurality of clients connected over the network to the servers, each of said clients being adapted to request processing to be carried out by one or more of the servers, the controller comprising:

20 means for obtaining at least one initial string of values representing configuration parameters for the distributed processing system defining which server is to be accessed by which client;

modelling means for modelling the distributed
25 processing system using the configuration parameters and information on the communications speed within the

distributed processing system to determine a performance value for the or each initial string of values;

generating means for repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position and changing the value of one or more of the values of the resulting string of values to generate a new string of values, wherein said modelling means is adapted to determine a performance value for the configuration parameters represented by said new string of values;

determining means for determining optimum configuration parameters represented as one of said initial or new string of values for which the performance value is the best; and

control means for configuring the distributed processing system to control which server is to be accessed by which client in accordance with the determined optimum configuration parameters.

36. A controller according to claim 35, wherein said generating means is adapted to consider a last value in a said string of values as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values

sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

5 37. A controller according to claim 35 or claim 36, wherein said obtaining means is adapted to randomly obtain values for said at least one initial string of values.

10 38. A controller according to any one of claims 35 to 37, wherein said obtaining means is adapted to obtain a plurality of said initial strings of values to form a population, and said generating means is adapted to perform a genetic algorithm on said population by
15 repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in a first said string of values,

(b) replacing a sequence of values of the same
20 length in a second said string of values at a random position, and

(c) mutating the value of one or more of the values of the resulting string of values to generate said new string of values.

25

39. A controller according to claim 38, wherein said generating means is adapted to repeatedly

(a) delete a proportion of the population for which the performance values are worst, and

(b) perform said repeated generation of new strings of values until the population is regenerated; and

5 to select said first and second strings of values randomly from the undeleted proportion of the population for which the performance values are best.

40. A controller according to claim 38, wherein said
10 generating means is adapted to repeatedly

(a) randomly select three said strings of values from said population,

(b) select two of the selected strings of values for which the performance values are worst, and

15 (c) replace the third of said selected strings of values with the generated new string of values.

41. A controller according to any one of claims 35 to 37, wherein said obtaining means is adapted to obtain a
20 single said initial string of values as a parent string of values and said generating means is adapted to repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in said parent string of
25 values,

(b) replacing a sequence of values of the same length in said parent string of values at a random position, and

(c) changing the value of one or more of the values
5 of the resulting string of values to generate said new string of values.

42. A controller according to claim 41, wherein said
10 generating means is adapted to repeatedly replace said parent string of values with said new string of values if the performance value associated with the model having said new string of parameters is best.

43. A controller according to claim 41, wherein said
15 generating means is adapted to repeatedly replace said parent string of values with said new string of values if the exponential of the difference in the performance values for said parent string of values and said new string of values divided by a factor dependent upon the
20 number of repetitions by said generating means is greater than a random number between 0 and 1.

44. A controller according to any one of claims 35 to
25 43, wherein said modelling means is adapted to determine said performance value in dependence upon the transaction times by said servers.

45. A controller according to claim 44 wherein said modelling means is adapted to determine said performance value as the longest transaction time by any said server.

5 46. A controller according to claim 44, wherein said modelling means is adapted to determine said performance value as the longest transaction time by any said data server and a proportion of the average of the transaction times for all said servers.

10

47. A controller according to any one of claims 35 to 46 including monitoring means for monitoring the transaction times of said servers, wherein said modelling means is adapted to use the output of said monitoring means to determine said performance values, and said control means is adapted to adaptively configure the distributed processing system.

15 48. A processor implemented method of determining optimum parameters of a model of a physical system, the method comprising:

obtaining at least one initial string of values representing the parameters of the model to be optimised;

determining a cost value associated with the model
25 having parameters represented by the or each string of values;

repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a
5 said string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values;

determining a cost value associated with the model having parameters represented by the new string of
10 values; and

determining the optimum parameters as one of said initial or new string of values for which the cost value is closest to a target.

15 49. A method according to claim 48, wherein in the generating step a last value in a said string of values is considered as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values
20 sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

50. A method according to claim 48 or claim 49, wherein
25 the values for said at least one initial string of values is randomly obtained.

51. A method according to any one of claims 48 to 50, wherein a plurality of said initial strings of values are obtained to form a population, and the repeated generating step comprises selecting a sequence of values
5 of random length starting at a random position in a first said string of values, replacing a sequence of values of the same length in a second said string of values, and changing the value of one or more of the values of the resulting string of values to generate said new string
10 of values.

52. A method according to claim 51, including repeatedly:

(a) deleting a proportion of the population for
15 which the cost value is furthest from said target;

(b) performing said repeated generating step until the population is regenerated; and

(c) selecting said first and second strings of values randomly from the proportion of the population for
20 which the cost values are closest to said target.

53. A method according to claim 51 including repeatedly:

(a) randomly selecting three said strings of values from said population

25 (b) selecting two of the selected strings of values for which the cost values are closest to said target as

said first and second strings of values for use in the generating step; and

(c) replacing the third of said selected string of values with the generated new string of values.

5

54. A method according to any one of claims 48 to 50, wherein a single said initial string of values is obtained as a parent string of values, and the repeated generating step comprises selecting a sequence of values
10 or random length starting at a random position in said parent string of values, replacing a sequence of values of the same length in said parent string of values at a random position, and changing the value of one or more of the values of the resulting string of values to
15 generate said new string of values.

55. A method according to claim 54, including repeatedly replacing said parent string of values with said new string of values if the cost value for said new string
20 of values is closer to said target.

56. A method according to claim 55, including repeatedly replacing said parent string of values with said new string of values if the exponential of the difference in
25 the cost values for said parent and new strings of values divided by as factor dependent upon the number of

repetitions of the generating step is greater than a randomly generated number between 0 and 1.

57. A method of controlling the configuration of a physical system, the method comprising:

obtaining at least one initial string of values representing configuration parameters of the physical system;

modelling the physical system using said configuration parameters to determine a cost value for the or each initial string of values;

repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position and changing the value of one or more of the values of the resulting string of values to generate a new string of values;

modelling the physical system using said configuration parameters to determine cost values for the new strings of values;

determining optimum configuration parameters represented as one of said initial or new string of values for which the cost value is closest to a target; and

controlling the configuration of said physical system in accordance with the determined optimum configuration parameters.

5 58. A method according to claim 57, wherein in the generating step a last value in a said string of values is considered as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values
10 sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

15 59. A method according to claim 57 or claim 58, wherein the values for said at least one initial string of values is randomly obtained.

20 60. A method according to any one of claims 57 to 59, wherein a plurality of said initial strings of values are obtained to form a population, and the repeated generating step comprises selecting a sequence of values of random length starting at a random position in a first said string of values, replacing a sequence of values of the same length in a second said string of values, and
25 changing the value of one or more of the values of the resulting string of values to generate said new string of values.

61. A method according to claim 60, including repeatedly:

(a) deleting a proportion of the population for which the cost value is furthest from said target;

5 (b) performing said repeated generating step until the population is regenerated; and

(c) selecting said first and second strings of values randomly from the proportion of the population for which the cost values are closest to said target.

10

62. A method according to claim 61 including repeatedly:

(a) randomly selecting three said strings of values from said population;

15 (b) selecting two of the selected strings of values for which the cost values are closest to said target as said first and second strings of values for use in the generating step; and

(c) replacing the third of said selected string of values with the generated new string of values.

20

63. A method according to any one of claims 57 to 59, wherein a single said initial string of values is obtained as a parent string of values, and the repeated generating step comprises selecting a sequence of values
25 of random length starting at a random position in said parent string of values, replacing a sequence of values of the same length in said parent string of values at a

random position, and changing the value of one or more of the values of the resulting string of values to generate said new string of values.

- 5 64. A method according to claim 63, including repeatedly replacing said parent string of values with said new string of values if the cost value for said new string of values is closer to said target.
- 10 65. A method according to claim 63, including repeatedly replacing said parent string of values with said new string of values if the exponential of the difference in the cost values for said parent and new strings of values divided by as factor dependent upon the number of
- 15 repetitions of the generating step is greater than a randomly generated number between 0 and 1.
- 20 66. A method according to any one of claims 57 to 65, including receiving operating parameters from said physical system, wherein the modelling step includes using the received operating parameters to model said physical system.
- 25 67. A method of configuring a distributed database comprising a plurality of data servers connected over a network, each data server holding any number of data units, and a plurality of clients connected over said

network to said data servers, each of said client being adapted to retrieve data units from and/or update data units at one or more of the data servers, the method comprising:

- 5 obtaining at least one initial string of values representing configuration parameters for the distributed database defining which data server is to be accessed by which client;
- 10 modelling the distributed database using the configuration parameters and information on communications within the distributed database to determine a performance value for the or each initial string of values;
- 15 repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position, and changing the value of one or more of the values of the resulting
- 20 string of values to generate a new string of values,
- modelling the distributed database using the configuration parameters and information on the passage of data within the distributed database to determine a performance value for the new strings of values;
- 25 determining optimum configuration parameters represented as one of said initial or new string of values for which the performance value is the best; and

configuring the distributed database to control which data server is to be accessed by which client in accordance with the determined optimum configuration parameters.

5

68. A method according to claim 67, wherein in the generating step a last value in a said string of values is considered as being segmented to a first value in said string of values such that said selected sequence of values can include said last and first values sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

15 69. A method according to claim 67 or claim 68, wherein the values for said at least one initial string of values is randomly obtained.

20 70. A method according to any one of claims 67 to 69, wherein a plurality of said initial strings of values are obtained to form a population, and the repeated generating step comprises selecting a sequence of values of random length starting at a random position in a first said string of values, replacing a sequence of values of the same length in a second said string of values, and
25 changing the value of one or more of the values of the

resulting string of values to generate said new string of values.

71. A method according to claim 70, including
5 repeatedly:

(a) deleting a proportion of the population for which the performance values are worst;

(b) performing said repeated generating step until the population is regenerated; and

10 (c) selecting said first and second strings of values randomly from the proportion of the population for which the performance values are best.

72. A method according to claim 70, including
15 repeatedly:

(a) randomly selecting three said strings of values from said population;

(b) selecting two of the selected strings of values for which the performance values are best; and

20 (c) replacing the third of said selected string of values with the generated new string of values.

73. A method according to any one of claims 67 to 69, wherein a single said initial string of values is
25 obtained as a parent string of values, and the repeated generating step comprises selecting a sequence of values of random length starting at a random position in said

parent string of values, replacing a sequence of values of the same length in said parent string of values at a random position, and changing the value of one or more of the values of the resulting string of values to
5 generate said new string of values.

74. A method according to claim 73, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values
10 if the performance value associated with the model having said new string of parameters is best.

75. A method according to claim 73, wherein said generating means is adapted to repeatedly replace said
15 parent string of values with said new string of values if the exponential of the difference in the performance values for said parent string of values and said new string of values divided by a factor dependent upon the number of repetitions by said generating means is greater
20 than a random number between 0 and 1.

76. A method according to any one of claims 67 to 75, wherein the step of configuring the distributed database comprises copying, moving and/or deleting data units at
25 data servers to change the distribution of data units across the data servers, and/or changing the data servers to be accessed by said clients using the determined

optimum configuration parameters to improve the system performance.

- 5 77. A method according to any one of claims 67 to 76, wherein said performance values is determined in dependence upon the time taken for one or more said clients to retrieve and/or update data units at one or more data servers.
- 10 78. A method according to claim 77, wherein the performance value is determined as the longest response time for any said client means to access any said data server.
- 15 79. A method according to claim 77, wherein said performance value is determined as the longest response time for any said data server and a proportion of the average of the response times for all said data servers.
- 20 80. A method according to claim 77, wherein said performance value is determined as a function of the rates at which data can be retrieved by said clients from said data servers, the rates at which data can be updated by said clients at said data servers, the contention
25 between said clients for accessing said data servers, and communication times between said clients and said data servers.

81. A method according to any one of claims 67 to 80 including the step of monitoring the passage of information over the network, wherein the performance value is determined in accordance with the monitoring, and the configuration of the distributed database is adaptively controlled.

82. A method of configuring a distributed processing system comprising a plurality of servers connected over a network, each server being capable of carrying out any number of processing operations, and a plurality of clients connected over said network to said servers, each of said client being adapted to request processing to be carried out by one or more of the servers, the method comprising:

obtaining at least one initial string of values representing configuration parameters for the distributed processing system defining which server is to be accessed by which client;

modelling the distributed processing system using the configuration parameters and information on communications within the distributed database to determine a performance value for the or each initial string of values;

repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values,

replacing a sequence of values of the same length in a said string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values,

5 modelling the distributed processing system using the configuration parameters and information on the passage of data within the distributed processing system to determine a performance value for the new strings of values;

10 determining optimum configuration parameters represented as one of said initial or new string of values for which the performance value is the best; and

 configuring the distributed processing system to control which server is to be accessed by which client
15 in accordance with the determined optimum configuration parameters.

83. A method according to claim 82, wherein in the generating step a last value in a said string of values
20 is considered as being segmented to a first value in said string of values such that said selected sequence of values can include said last and first values sequentially in a said string of values and the sequence of values to be replaced can include said last and first
25 values sequentially in a said string of values.

84. A method according to claim 82 or claim 83, wherein the values for said at least one initial string of values is randomly obtained.

5 85. A method according to any one of claims 82 to 84, wherein a plurality of said initial strings of values are obtained to form a population, and the repeated
10 generating step comprises selecting a sequence of values of random length starting at a random position in a first said string of values, replacing a sequence of values of the same length in a second said string of values, and changing the value of one or more of the values of the resulting string of values to generate said new string of values.

15

86. A method according to claim 85, including repeatedly:

(a) deleting a proportion of the population for which the performance values are worst;

20 (b) performing said repeated generating step until the population is regenerated; and

(c) selecting said first and second strings of values randomly from the proportion of the population for which the performance values are best.

25

87. A method according to claim 85, including repeatedly:

(a) randomly selecting three said strings of values from said population;

(b) selecting two of the selected strings of values for which the performance values are best; and

5 (c) replacing the third of said selected string of values with the generated new string of values.

88. A method according to any one of claims 82 to 84, wherein a single said initial string of values is
10 obtained as a parent string of values, and the repeated generating step comprises selecting a sequence of values of random length starting at a random position in said parent string of values, replacing a sequence of values of the same length in said parent string of values at a
15 random position, and changing the value of one or more of the values of the resulting string of values to generate said new string of values.

89. A method according to claim 88, wherein said
20 generating means is adapted to repeatedly replace said parent string of values with said new string of values if the performance value associated with the model having said new string of parameters is best.

25 90. A method according to claim 88, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values

if the exponential of the difference in the performance values for said parent string of values and said new string of values divided by a factor dependent upon the number of repetitions by said generating means is greater
5 than a random number between 0 and 1.

91. A method according to any one of claims 82 to 90, wherein said performance values is determined in dependence upon the transaction time taken for said
10 servers.

92. A method according to claim 91, wherein the performance value is determined as the longest transaction time for said data server.
15

93. A method according to claim 91, wherein said performance value is determined as the longest transaction time for any said server and a proportion of the average of the transaction times for all said
20 servers.

94. A method according to any one of claims 82 to 93 including the step of monitoring the transaction times of said servers, wherein the performance value is
25 determined in accordance with the monitoring, and the configuration of the distributed processing system is adaptively controlled.

95. A storage medium storing processor implementable instructions for controlling a processor to carry out the method of any one of claims 48 to 94.

- 5 96. An electronic signal carrying computer code for controlling a processor to carry out the method of any one of claims 48 to 94.

09/743520 5640



PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

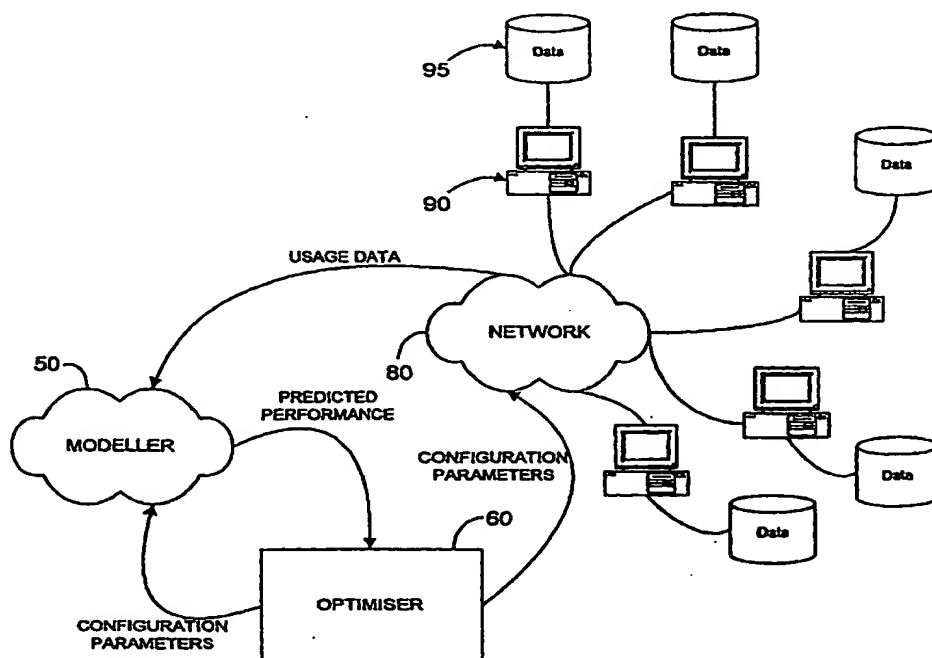
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : G06F 17/30, 15/18, 11/34, H04L 12/56		A1	(11) International Publication Number: WO 00/08569
			(43) International Publication Date: 17 February 2000 (17.02.00)
(21) International Application Number: PCT/GB99/02449		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 27 July 1999 (27.07.99)		<p>Published With international search report.</p>	
(30) Priority Data: 9817051.7 5 August 1998 (05.08.98) GB			
(71) Applicant (for all designated States except US): BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).			
(72) Inventor; and (75) Inventor/Applicant (for US only): OATES, Martin, John [GB/GB]; 11 Naughton Gardens, Stowmarket, Suffolk IP14 2RA (GB).			
(74) Agent: ROBINSON, Simon, Benjamin; BT Group Legal Services, Intellectual Property Dept., Holborn Centre, 8th Floor, 120 Holborn, London EC1N 2TE (GB).			

(54) Title: DATA PROCESSING APPARATUS AND METHOD FOR OPTIMISING CONFIGURATION PARAMETERS OF A PHYSICAL SYSTEM

(57) Abstract

A technique for determining optimum parameters of a model of a physical system is disclosed in which at least one initial string of values representing the parameters of the model to be optimised is obtained. A cost value associated with the model having parameters represented by the each string of values is determined. A new string of values is repeatedly generated by selecting a sequence of values of random length, starting at a random position in a string of values, replacing a sequence of values of the same length in a string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values. The cost value associated with the model having parameters represented by the new string of values is determined and the optimum parameters are determined as one of the initial or new string of values for which the cost value associated with the model having the optimum parameters is closest to a target.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

DATA PROCESSING APPARATUS AND METHOD FOR
OPTIMISING CONFIGURATION PARAMETERS
OF A PHYSICAL SYSTEM

5 The present invention generally relates to the data processing method and apparatus for determining optimum parameters of the model of a physical system and for controlling the configuration of the physical system.

 The generation of a model of a physical system is
10 used for a wide variety of applications including learning more about the behaviour of the system and controlling the parameters of the system. The model of the physical system can be used to try out parameters in order to determine optimum parameters which can then be
15 used for controlling the physical system. In this way the effect of the parameters on the physical system can be predicted so that only optimum parameters are chosen for use in the control of the physical system.

 One such physical system for which configuration
20 parameters are required is a distributed database. Because the use of corporate intranets continues to rise, the management of efficient access to data is becoming a key issue. Large information systems, often accessed on a global basis, are increasingly being provided by
25 distributed means with use of "mirror sites" to ease congestion and improve local access. The configuration of the information system defining the location of the

information and the servers to which clients are to direct queries is important in order to provide a system which has a high performance as perceived by the clients. Further, where access is truly global, static design solutions are not ideal, as the largest source of load on the system moves from geographic area to geographic area at different times of the day. This tends to create congestion that is "localised" in the region of demand. Localised mirroring in one region only will simply shift the contention onto the global communications networks when demand shifts to other regions. Mirroring in all regions is not only costly, but also leads to over duplication of data, and greater problems with integrity and the need to perform multiple updates of the data.

Where the data is required to be both updated as well as read, the administration of these information systems can become significant and labour intensive. In a recent paper by MJ Oates *et al* entitled "Investigating Evolutionary Approaches for Self-Adaption in Large Distributed Databases" (Proceedings of the 1998 IEEE International Conference on Evolutionary Computation) and in a paper by G Bilchev *et al* entitled "Comparing Evolutionary Algorithms and Greedy Heuristics for Adaption Problems" (Proceedings of the 1998 IEEE International Conference on Evolutionary Computation) it has been shown that autonomous management of distributed information systems is feasible and in these publications

the use of various types of algorithms for the adaption of distributed databases has been considered.

The present invention provides an improved method of apparatus for determining optimum parameters of the model of a physical system wherein at least one initial string of values representing the parameters of the model to be optimised is obtained; a cost value associated with the model having parameters represented by the reached string of values is determined; a new string of values is repeatedly generated by a) selecting a sequence of values of random length starting at a random position in a string of values, b) replacing the sequence of values of the same length in a string of values at a random position and c) changing the value of one more of the values of the resulting string of values to generate a new string of values; a cost value associated with the model having parameters represented by the new string of values is determined; and the optimum parameters are determined as one of the initial or new string of values for which the cost value associated with the model having the optimum parameters is closest to a target, such as a maximum or a minimum.

Using this technique new strings of values representing new parameters of a model are produced which can be applied to the model to see if a cost value is closer to a target e.g. a minimum or a maximum is

produced. The technique thus provides a novel search technique to search for the optimum parameters.

This technique is particularly applicable for the control of the configuration of a physical system wherein
5 the model models the physical system and the new parameters generated represent new untried configurations of the physical system. The cost value output from the model can be used to determine whether the new configurations are an improvement or not. Once the
10 optimum parameters have been determined they can be used to control the configuration of the physical system.

The present invention is particularly suited to the problem of configuring a distributed database comprising a plurality of data servers connected over a network,
15 each data server holding any number of data units, and a plurality of clients connected over the network to the data servers, each of said clients being adapted to retrieve data units from and/or update data units at one or more of the data servers. When the invention is
20 applied to the configuration of a distributed database, the configuration parameters define which data server is to be accessed by which client and the distributed database is modelled using the configuration parameters and information on the passage of data within the
25 distributed database to determine a performance value for each of the strings of values. Which ever of the strings of values have the best performance value is then used

for configuring the distributed database to control which data server is to be accessed by which client. The performance value is representative of the performance perceived by one or more of the clients. In this aspect of the present invention a client comprises any application running on a computer connected to the network which requires access to the data units stored on the data servers. Where the performance value represents the transaction time for one or more clients accessing a data unit on a data server, the best performance value is a minimum and thus the technique of this embodiment of the present invention will attempt to minimise the performance value i.e. the transaction time.

Although the present invention is particularly applicable to optimising the configuration of a distributed database, it can be applied to any physical system or model wherein parameters of the model can be provided as a string of values. The values can either comprise actual numerical values of the model or can comprise numerical values representing features of the model i.e. they comprise or represent symbols.

The present invention can be applied generally to the configuration of a distributed processing system, of which a distributed data base is but one example, wherein the configuration parameters determine which nodes (servers) in the network are to be used for processing requested by applications (clients).

In a preferred embodiment of the present invention the string of values is considered to wrap round to form a continuous string such that the last and first values of the string are sequential. In this way, when a
5 sequence of values at random length is selected in a string of values the sequence can include the last and first values sequentially. Similarly, when a sequence of values is replaced in a string of values this can include the sequential last and first values.

10 By selecting a sequence of values and overlaying it the intention is that good segments of the string of values in one part of the sting may also work well as a building block elsewhere in the string. This technique allows the values which may have become "extinct" in
15 positions in the population of strings to be "reinjecte" thereby enhancing the searching technique. The technique for generating the initial parent strings of values does not appear to be important. In one embodiment these are generated randomly.

20 The present invention is applicable to evolutionary computational search techniques and other conventional search techniques. For example, the technique of the present invention can be applied to a genetic algorithm wherein a population of strings of values is created and
25 new strings of values are created from two parents. The technique of the present invention when applied to genetic algorithms is a cross-over technique wherein a

sequence of values of random length is selected at a random start position in a first string of values and this is used to replace a sequence of values at the same length in a second string of values at a random position.

5 One or more of the values of the resulting string of values is then mutated to generate the new string of values. This basic technique can be used in any sort of genetic algorithm such as a Breeder genetic algorithm or a Tournament genetic algorithm.

10 As mentioned above, the present invention is applicable to non-evolutionary search techniques in which a single string of values is used as a parent from which a new string of values is generated. In this method the sequence of values of random length is selected starting
15 at a random position in the parent string of values and this is used to replace a sequence of values in the same length in the parent string of values at a random position. One or more of the values of the resulting string of values is then changed to generate the new
20 string of values. This basic technique can be used in a Hill Climbing search technique or in a Simulated Annealing search technique for example.

Since the strings of values comprise a number of values, each of which is a value for a particular
25 parameter, the strings of values can be termed solution vectors since they can be represented as n dimensional

vectors and the "best" one will comprise the solution to the problem of finding the best configuration parameters.

Embodiments to the present invention will now be
5 described with reference to the accompanying drawings, in which:

Figure 1 is a schematic illustration of the use of the technique of the present invention to determine optimum parameters of a model in the physical system,

10 Figure 2 is a schematic diagram of use of the technique of Figure 1 to control the configuration of a physical system,

Figure 3 is a schematic diagram illustrating the use of the technique to control the configuration of a distributed database in accordance with one embodiment
15 of the present invention,

Figure 4 is a schematic diagram of a distributed database in accordance with one embodiment of the present invention,

20 Figure 5 is a schematic diagram of a distributed database in accordance with another embodiment of the present invention,

Figure 6 is a schematic diagram of a distributed database in accordance with a further embodiment of the
25 present invention,

Figure 7a is a flow diagram illustrating a retrieval operation,

Figure 7b is a flow diagram illustrating an update operation,

Figure 8 is a flow diagram illustrating steps in adapting the distributed database,

5 Figure 9 is a schematic diagram of a distributed database in accordance with a first scenario,

Figure 10 is a schematic illustration of a solution vector,

10 Figure 11 is a flow diagram of the steps in a first "basic" method for using a model to determine a cost value,

Figure 12 is a flow diagram of a Breeder genetic algorithm for determining the lowest cost value,

15 Figure 13 is a flow diagram of a single three way Tournament genetic algorithm for determining the lowest cost value,

Figure 14 is a flow diagram of the steps for generating a new solution vector in accordance with an embodiment of the present invention,

20 Figure 15 is a schematic illustration of the process for generating a new solution vector in accordance with the flow diagram of Figure 14,

Figure 16 is a schematic illustration of a distributed database of a second scenario (B),

25 Figure 17 is a flow diagram illustrating the steps of a second method termed "just used" for using the model to determine a cost value,

Figure 18 is a flow diagram illustrating the steps of a third method termed "plus average" for using the model to determine a cost value,

Figure 19 is a flow diagram illustrating the steps of a fourth method termed "plus used" for using the model to determine a cost value,

Figure 20 is a flow diagram illustrating the steps of a fifth method termed "just used" for using the model to determine a cost value,

Figure 21 is a flow diagram illustrating the steps of a sixth method termed "plus all" for using the model to determine a cost value,

Figure 22 is a flow diagram illustrating the steps of a seventh method termed "plus 10%" for using the model to determine a cost value,

Figure 23 is a flow diagram illustrating a Hill Climbing algorithm,

Figure 24 is a flow diagram illustrating a Simulated Annealing algorithm,

Figure 25 is a flow diagram illustrating the steps of one embodiment of the present invention for generating a new solution vector,

Figure 26 is an illustration of the technique of Figure 25 for generating a new solution vector,

Figure 27 is an illustration of the results for five different optimisation techniques for scenario A using the basic, least worst server model,

Figure 28 is an illustration of the results for the five different optimisation techniques for scenario A using the "plus all" model,

Figure 29 is an illustration of the results for the
5 five different techniques for scenario B using the basic, least worst model, and

Figure 30 is an illustration of the results for the five different techniques for scenario B using the "plus all" model.

10

Figure 1 is a schematic illustration of the principles of one aspect of the present invention. A modeller 50 operates to generate a model of a physical system. In order to generate the model the modeller 50
15 receives configuration parameters from an optimiser 60. Thus, the optimiser 60 controls the configuration of the model generated by the modeller 50. As a result of the generated model, the modeller 50 outputs a cost value to the optimiser 60. The cost value can represent any
20 function of the model which it is desired to, for example, reduce or increase. For example, in the design of a system it may be desirable for the cost value to represent the energy consumption of the system which is to be reduced or the performance of the system which is
25 to be improved (increased). The optimiser 60 operates a form of a search algorithm in order to generate a string of values representing the configuration

parameters. The string of values is termed a solution vector since it can be considered to be an n-dimensional vector having values representing the configuration parameters. The optimiser 60 adaptively changes the configuration parameters in order to optimise the cost value i.e. either reduce it or increase it. The search for a solution vector which will provide the optimum cost value involves the copying of a segment of a solution vector of random length and starting at a random position as an overlay for a segment of the solution vector of the same length at another random position. One or more of the values of the solution vector is then chosen randomly and changed. This process generates a new solution vector representing a new set of configuration parameters used by the modeller 50. The generation of the new configuration parameters using this technique provides an improved method of searching for the optimum cost value.

Figure 2 is a schematic illustration of a second aspect of the present invention wherein the technique of Figure 1 is applied to the control of the configuration of a physical system 70. The modeller 50 and optimiser 60 operate as described hereinabove with reference to Figure 1. Once the optimiser has determined an optimum cost value, the configuration parameters corresponding to the optimum cost value are output to the physical system 70 in order that the system be optimally

configured. In order for the modeller 50 to accurately model the physical system 70, the physical system outputs operating parameters to the modeller 50. In this aspect of the present invention the use of the modeller 50 enables the prediction of the outcome of changes in the configuration of the physical system before the changes are actually made to the physical system 70. In other words, no non-optimum changes need be made to the configuration of the physical system 70 in order to try to achieve an optimum configuration. The determination of the optimum configuration is predicted using the model.

The operating parameters output from the physical system 70 to the modeller 50 can comprise any measured data which is required for the accurate modelling of the physical system by the modeller 50. Such data can be measured from the physical system.

Figure 3 is a schematic illustration of an aspect of the present invention wherein a distributed database is optimally configured. This system is as described with reference to Figure 2 except that the physical system is a distributed database comprising data servers 95 and terminals 90 operating clients requiring access to data on the databases 95 and interconnected via a network 80. The modeller 50 receives usage data indicating the pattern of usage of the data in the databases 95 by the client's in the terminals 90. The

modeller 50 also requires information on communication speeds in order to accurately model the distributed database. The optimiser 60 will search for an optimum solution vector representing configuration parameters for the distributed database which indicate the optimum usage pattern i.e. which databases 95 the clients on the terminals 90 should look to for retrieval of data. if their usage patterns perform a repeated cycle or are predictable, projected usage data can be fed into the model instead of retrieved from the network 80 thereby allowing the system to generate the new configuration parameters based on anticipated work load. The data units held in the databases 95 can comprise any form of information such as text, audio, video or images, or any other form of database information.

The distributed database can be provided over a local area network such as an ethernet or a corporate intranet, or over a wide area network such as the internet.

The distributed database can either be a homogenous database wherein all of the data units are of the same format, or it can comprise a heterogenous database wherein data units are stored in different formats. In a heterogenous database it will be necessary to translate between the different formats of the data units in order for data to be passed between databases e.g. during data "mirroring" operations, and during retrieval and update

operations if the format of the data operated on by the client is different to the format of data in the database.

The reconfiguration of the database based on the configuration parameters results in a client being directed to a particular data server for data. When the configuration is changed the client is directed to a new data server, if the data server does not contain the required data it is necessary for the data to be copied from a data server containing that data. Thus, the reconfiguration of the distributed database comprises a mix of simply redirection of queries from a client to a database server and the copying of data between data servers. Since the copying of data between data servers requires a high band width over the network, this operation is minimised where possible. One scenario, where the copying of data between data servers takes place is when a data server becomes heavily used by clients. The "mirroring" of the data to another data server will spread the load between the data servers thereby enhancing the performance of the distributed database.

The adaption of the configuration of the distributed database can either take place continuously during the operation of the distributed database, at a period when the use of the distributed databases are low, or at predetermined times.

The modeller 50 and optimiser 60 are preferably implemented as a computer program operating on a computer. In order to reduce the amount of processing that is carried out, the calculation of the optimum configuration parameters can be triggered to occur only in response to a suitable criteria such as the failure of a data server, a network link failure or change, or when the time taken by one or more of the clients to retrieve and/or update data units at one or more of the data servers exceeds a threshold. This avoids the adaption process continuously hunting for an optimum solution when the distributed database may already be optimumly configured or near optimumly configured. Alternatively, the adaption algorithm can operate continuously but the configuration determined may only be applied if the predicted improvement in performance exceeds a threshold. Since the overheads associated with reconfiguration can be significant, a balance must be struck between accepting sub-optimal performance and the frequency of reconfiguration.

Figure 4 illustrates a distributed database in one embodiment of the present invention wherein the first database 1 (DB1) and the second database 2 (DB2) are connected to a network 3 and thus form nodes in the network. The first database 1 is a relational database and includes a relational database management system (RDBMS) 4. The second database 2 is an object oriented

database and includes an object oriented database management system (OODBMS) 5. The database management systems 4 and 5 manage the respective database e.g. control access, security, backups etc. and operate to
5 interface the databases 1 and 2 to the network 3.

A first network station 6 is connected to the network 3 and operates a first application APP1 which is interfaced to the network 3 by an access manager AM1. A second work station 7 is also connected to the network
10 3 and runs a second application APP2 which is interfaced to the network 3 via an application manager AM2. The applications APP1 and APP2 include performance files 8 and 9 for the storage of the monitored performance during retrieval and/or updating of data to the databases 1 and
15 2. The access managers AM1 and AM2 include respective usage files 10 and 11 for storing information on the usage of databases 1 and 2. Also the access managers AM1 and AM2 include translators 12 and 13 for translating queries and data between database formats when necessary.

20 The database network also includes a controller 14 which operates an adaptor application 15 for carrying out the adaption algorithm for the control of the configuration of the database. The adaptor application 15 is interfaced with the network 3 via an adaptor
25 application access manager AM3. The access manager AM3 includes a translator 16 for translating data between different database formats where necessary and a parser

for parsing generic messages to move, copy and/or delete data. The controller 14 also contains tables 17 which contain data identifying the location of data files in the databases, information identifying the type of format
5 of the database and the location, and information identifying the database containing information which may be accessed by each application APP1 and APP2.

Although in Figure 4 only two work stations 6 and 7 and two databases 1 and 2 are illustrated, any number
10 can be provided over the network 3. Also only work stations 5 and 6 are illustrated as only operating single respective applications APP1 and APP2, they can, in a multitask environment, be operating more than one application which can retrieve and/or update data i.e.
15 the work stations 5 and 6 can host more than one client. The respective access manager AM1 and AM2 identifies which of the applications in a respective work station 6 and 7 has requested or sent data and it is routed accordingly.

20 Within Figure 4 the adaptor application 15 and the table 17 are illustrated as being provided in a separate controller arrangement 14. However, these can be distributed throughout the work stations 6 and 7 as is illustrated in Figures 5 and 6.

25 Figure 5 illustrates an embodiment wherein the work stations 6 and 7 include tables 17a and 17b which are copies of tables 17. This arrangement is advantageous

since the access managers AM1 and AM2 can access the tables directly without incurring a network delay.

Figure 6 illustrates a further embodiment wherein the work stations 6 and 7 also run local adaptor applications 15a and 15b and thus the adaption algorithm is carried out in a distributed manner. Local parsers 20a and 20b are provided in the respective access managers AM1 and AM2 for parsing generic messages from the respective adaptor applications 15a and 15b. This distribution of the adaptor application will benefit simply from parallel processing, or the adaptor application 15a and 15b could operate on a localised basis to control distribution of data locally i.e. only at a few of the databases.

Methods of adaptively controlling the configuration of the distributed databases will now be described.

Figure 7a illustrates the steps carried out when data is requested from a database by an application on a work station. In step S1 the application generates a query for data which includes the data identification (ID). In step S2 the access manager responds to the query and looks in the tables to determine the application location, the database location for the applications and the identity and type of database containing data. In step S3, based on the information from the tables, it is determined whether there is a difference in format between the format of the database and the format of the

application. In step S4 if there is a difference in format the query is translated and in step S5 the data is obtained from the database. In step S6 the access module records the database usage in the usage file. If
5 the format of the database is different to the format of the application, in step S7 it is determined whether the translation of the data is necessary and in step S8 the translation of the data takes place if necessary. In step S9 the data is returned to the application by the access
10 manager. In step S10 the application records, in a performance file, how long the transaction took to complete which is an indication of the performance of the retrieval operation.

In step S11 the adaptor application reads the
15 performance and usage files and in step S12 the adaptor application can determine whether an adaption of the distributed database is necessary. This adaption can only take place off line, continuously in real time, only at predetermined times, or only when network traffic will
20 permit the transference of data over the network between databases.

The usage file used in step S6 contains an identification of the data record retrieved from the database, an identification of the database from which
25 it was retrieved, an identification of the application which retrieved the data record, and the date and time of retrieval. The usage file can either comprise a

single file for databases, or it can comprise a single file for each of the databases. This is used for determining an optimum configuration for the distributed database.

5 The performance files comprise a file for each application. The performance files contain data identifying the retrieved record data, the application identification, the query identification, and the transaction duration. These are used for determining
10 whether to determine an optimum configuration for the distributed database.

 In the embodiments described hereinabove the tables 17, 17a and 17b can comprise a number of different tables containing information used for identifying the location
15 of data in the distributed database. The system configuration table can contain data which identifies the database, the database type e.g. a relation database or an object oriented database, and the connect string used to connect to the database over the network. The
20 advantage of using such a table is that when a new database is added to the distributed database network, its details can simply be added to this file and the algorithm, as will be described hereinafter, can adapt the network to efficiently accommodate it.

25 The tables can also include a data allocation for each work station which records where an application on a network station should look for each data record. In

these embodiments of the present invention these tables specify where the information can be retrieved for an application.

The tables can also include a data configuration
5 table which lists all of the locations of each type of data record.

The table can further include an application
configuration table which is used by the access manager
so that it can determine where each application is
10 running and so which data allocation to use. The table
thus containing the application identification and the
identification of the work station on which it is
running.

Figure 7b is a flow diagram of steps carried out for
15 an update. In step S100 the application generates an
update query and identifies the data type which is to be
updated. The access manager receives the query and data
from the application in step S101 and looks in the tables
to determine the application location, the database
20 locations containing the data type and the identity of
the databases containing the type of data. In step S102
it is determined whether it is necessary to translate the
update query and the data from one database to another
depending upon the database format of the data generated
25 by the application and the data format of the type
database to be updated. If the translation between the
database format is required, in step S103 the update

query and data are translated. In step S104 the data is then transferred to the databases containing the data type to update the data therein. In step S105 the access manager records the database usage in the usage file and
5 in step S106 the database management system returns a configuration of the update of the application. The application then records the performance in the performance file in step S107 and in step S108 the adaptor application reads the performance and usage file.
10 If the adaptor application determine that the performance of the distributed database can be improved, in step S109 the adaptor application adapts the distributed database.

Thus in accordance with the update method, all of the database types which corresponds to the type of data
15 generated by the application are updated for subsequent access by any other application in the distributed database.

The adaptor algorithm in accordance with one embodiment of the present invention will now be
20 described.

Applications access data either on a local server or, via a communications link, on a remote server. The best location of data depends upon the concentrations of demand and the relative performance cost of each server
25 and the associated communications costs. An optimum configuration may require data replication so that the overall demand for given data can be shared across

multiple servers and so reduce load and performance costs for a given server. Whilst the distribution of data across multiple servers is desirable and advantageous for data retrieval, this poses difficulties for updating data since data must be updated at all servers. The algorithm aims to compute an improved cost solution to cause a reconfiguration of the distributed database when overall performance figures have exceeded an undesirable threshold. Other factors can trigger the adaption such as the failure or anticipated shut down of one or more communication links or databases, or communication costs for a communication link exceeding a threshold.

The following algorithm, in accordance with one embodiment of the present invention, computes costs into in terms of response times i.e. the time between initiating the request for data and the data being received.

The steps of the adaptation method of one embodiment of the present invention will now be described with reference to Figure 8.

In step S20 the adaption algorithm for the adaptor application is triggered into operation by a threshold decrease in performance being exceeded. In step S21 the adaptor application calculates the optimum distribution of data of the network and the optimum application to data links taking into consideration the model parameters as will be described hereinafter. In step S22 the

adaptor application determines whether there are differences in the optimum distribution data and the current distribution of data and generates messages to copy/move and/or delete data at database nodes in the network and to update links either define the nodes which are to receive queries from the applications i.e. where an application is to receive data. The adaptor application access manager receives the messages in step S23 and parses each message. In step S24 the adaptor application access manager uses the tables, in particular, the configuration table, to identify the database locations and types and send instructions to database drivers.

In step S25 the database drivers copy, move and/or delete data in accordance with the instructions and the adaptor application access manager translates data if the data is copied or moved from one database type to another. When all the data has been moved, copied and/or deleted, in step S26 the adaptor application access manager informs the adaptor application that all instructions have been carried out. In step S27 the adaptor application updates the tables and indicates any new locations of the data records in the network (i.e. updates the data configuration file) and updates the allocation file to specify where each application can be find appropriate data in databases in the distributed databases.

If data has to be moved from one database to another in order to increase performance, instead of moving the data, the data could be copied across one database to the other and subsequently deleted from the original database. This operation ensures that a copy of the data is always available for access. If the data is simply moved, the data may not be available during transit from one database to another. Further, if the movement of data is divided into a two stage copy and delete process, not all of the allocation tables need to be updated simultaneously. Before the deletion of data, some applications can still access data from the old locations until their data allocation tables have been updated to indicate the new locations. Thus dividing the move operation into a copy and delete two stage operation may have advantages in a system which operates a real time adaptive algorithm.

Although the adaption method of Figure 8 has been described as including the moving of data, to optimise the database configuration it is not always necessary to move data between servers. It may be possible to improve performance simply by changing the servers to which queries are sent by one or more of the applications. For example, where data is present for an application on two servers, and the one currently receiving the queries for retrieving data is under a heavy load, the adaption method can simply cause the application query to be

directed to the other server which is less heavily loaded.

Specific adaption algorithms for determining a data distribution and the data servers to which application queries are sent and which can be used to adapt the configuration of the distributed database to improve performance (access time) will now be described.

The methods use tables of figures which represent possible client server connections and related transaction rates and performance figures. These are modelling parameters which represent a model of the distributed database. The performance figures represent costs that are likely to be incurred for a given configuration and are derived from measured values for similar transaction types to the ones being considered.

A model of the distributed database used in the following embodiments includes as parameters forms of both data servers and communications networks so that the performance of the system seen by individual client applications can be estimated for the current load conditions over a range of different access and server configurations. The choice of configuration is determined by an optimisation algorithm which produces solution vectors defining, for each client, which server that client should currently connect to for read access (update access must be made to all servers maintaining a copy of the database).

The algorithm derives costs from retrieval and update rates which can be obtained from the usage files in the embodiments described hereinabove and the degree of contention between retrieval and update transactions, as well as contention between clients on the same server node. In the algorithm it is assumed that the server and communications performance can be based on an M/M/1 queue. This is based on random transaction arrival rates, exponential service times and the assumption of a single server per client. The principal equations used in the algorithms are:

$$\text{Response time} = 1 / ((1/BTT) - (1/TIT)) \dots \dots \dots (1)$$

where: BTT is the Base Transaction Time, and
TIT is the Transaction Interarrival Time

$$\text{Overall Transaction Rate (OTR)} = ((\text{sum of Retrieval and Update Transaction Rates} - \text{Maximum Transaction Rate}) \times \text{contention value}) + \text{Maximum Transaction Rate} \dots \dots (2)$$

Thus if the contention value equals 100%, then Overall Transaction Rate equals sum of Transaction Rates.

If the contention value equals zero, the Overall Transaction Rate equals the Maximum Transaction Rate.

For all the contention values, the Overall Transaction Rate will be between the two extremes.

Figure 9 is a schematic illustration of a first database scenario (A) wherein the distributed database comprises 10 databases 100b (A to J) and 10 clients 100a (1 to 10) connected over a network 200 wherein the communication speeds over the network are uniform. Each client 100a resides with a respective database 100b at a node 100 of the network 200. Thus each client can access its respective database without incurring a network delay.

The configuration of the distributed database can be described as a solution vector an example of which is illustrated in Figure 10. The solution vector comprises a string of numbers indexed by client and containing an identification of the server to which each client should access to read data. In this illustrated example it can be seen that a server may be used by more than one client e.g. clients 1 and 8 access server B and clients 3 and 9 access server A.

Two genetic algorithm methods will now be described for determining the optimum solution vector for the scenario illustrated in Figure 9.

Tables 1, 2 and 3 illustrate the modelling parameters used in these two genetic algorithm methods. The algorithm operates for a distributed database as illustrated in Figure 9 i.e. 10 clients and 10 servers and a server contention overlap of 10% is assumed.

TABLE 1

Server Base Transaction Times (BTT) per node (in milliseconds) for database scenario A:

<u>SERVER</u>	<u>BTT</u>
1	>3000
2	>2800
4	>3100
4	>2900
5	>4000
6	>5000
7	>3500
8	>4000
9	>4500
10	>3750

TABLE 2

Client Application rates per node: Retrieval Rate (RR), Update Rate (UR), Overlap % (in events per second) for database scenario A

<u>CLIENT</u>	<u>RR</u>	<u>UR</u>	<u>OVERLAP</u>
4	>0.1	0.01	20
2	>0.2	0.01	20
3	>0.1	0.01	20
4	>0.1	0.01	20
5	>0.15	0.01	20
8	>0.1	0.01	20
8	>0.1	0.01	20
8	>0.15	0.02	25
9	>0.1	0.01	15
10	>0.1	0.02	15

TABLE 3

Base Comms Table (time is milliseconds) for database scenario A

Server/ Client	1	2	3	4	5	6	7	8	9	10
1	400	1000	2000	1900	1500	1200	1600	1700	1850	1800
2	1000	400	2000	2000	1800	1900	1750	1650	1800	1900
3	1000	800	400	2000	1750	1600	1670	1760	2000	1900
4	800	1200	2000	400	1850	1900	1950	1850	1900	2000
5	900	1100	1900	1700	400	1800	1600	1760	1640	1860
6	950	1050	1950	1800	1600	400	1550	1750	1750	1800
7	1000	1100	1900	1500	1650	1800	400	1700	1700	1800
8	900	950	1800	1900	1800	1500	1700	400	1900	1950
9	850	1050	1850	1950	1750	1700	1650	1700	400	1900
10	850	1000	1800	1900	1800	1750	1700	1650	1870	400

Figure 11 illustrates a method of a first embodiment of the present invention for modelling the distributed database using a proposed solution vector in order to generate the cost value which, for this embodiment, comprises a worst retrieval rate for the database.

In step S40 the effective transaction rate (ETR) is calculated for each client by combining the retrieval and update rates, using the percentage overlap and equation 2 given hereinabove. In step S41 a 2-dimensional array L is populated with the update rate calculated for that client and indexed by client and server number. Where a server is to be used by a particular client as indicated in the input current proposed solution from the genetic algorithm (step S43), in step S42 the

corresponding array entry in L is replaced with that client's effective transaction rate. A 2-dimensional array C is then created in step S44 to hold effective communication overhead delays. For each client/server
5 interaction rate, now defined in L, the corresponding entry in C is calculated using equation 1 given hereinabove. In step S45, for each server, the client entries in L are used to calculate aggregate loading from all clients of that server using equation 2 given
10 hereinabove. In step S46 this is then converted to an average response time using equation 1 given hereinabove and this is stored in a 1-dimensional array T indexed by server. Any calculated infinite response times are replaced with a suitably large response time. In step
15 S47, for each server, the client with the worst communications overhead delay is found from C and this is added to the relevant entry in T. In step S48 the largest value T is then output and this represents the worst performance of this database configuration.

20 A first algorithm for generating proposed solution vectors will now be described herein after with reference to Figure 12. This flow diagram illustrates the steps of a Breeder genetic algorithm.

In step S50 an initial random population P is
25 created using a non-binary representation. Each gene position corresponds to a client node and the allele indicates which server that client is to be used for

retrieval access. The maximum number of generations G to be allowed is calculated in step S51 from the following equation:

5 $G = 5000 / ((\text{population size}/2)+1) \dots \dots \dots (3)$

 In step S52 all the members of the population are then evaluated using the method of Figure 11. In step S53 g is set to 0. In step S54 the current generation
10 number g is incremented by 1 and a loop in the algorithm is entered. All of the numbers of the population are sorted in step S55 based on the evaluation result such that the lowest result is sorted to the top i.e. is the best. The bottom half of the population is then deleted
15 in step S56 and thus the current population p is set to equal half of the total population P . In step S57 the current population p is incremented by 1 and in step S58 two members from the top half of the population are chosen at random and a new number is generated using the
20 technique which will be described hereinafter with reference to Figures 14 and 15. In step S59 using uniformly distributed allele replacement each gene is mutated in the new member based on the defined percentage chance of mutation. In step S60 the new member is
25 evaluated using the procedure of Figure 11 and this is added to the bottom of the population list. In step S61 it is then determined whether the original population

size had been restored i.e. $p = P$ and if not the process returns to step S57. If the original population size P has been restored the process proceeds to step S62 whereupon it is determined whether the maximum number of generations G has been reached i.e. $g = G$. If $g \neq G$ the process returns to step S54. If $g = G$ the process proceeds to step S63 where all the members of the population are sorted based on the evaluation results from the lowest result and best. In step S64 the member of the population with the lowest evaluation result is entered. This can then be used for determining the configuration of the distributed database.

Figure 13 illustrates an alternative method which implements a single three way Tournament genetic algorithm. In step S70 an initial random population is created using a non-binary representation. Each gene position corresponds to a client node and the allele indicates which server that client is to use for retrieval access. In step S71 the maximum number of additional evaluations A to be allowed is calculated using the following equation:

$$A = 5000 - \text{population size} \dots \dots \dots (4)$$

In step S72 all the members of the population are evaluated using the method of Figure 11. In step S73 the current evaluation a is set to 0 and in step S74 the

current evaluation A is incremented to enter a loop in the algorithm. Three members of the population are then chosen at random in step S75 and in step S76 these are sorted into BEST, SECOND and WORST. A new member is created in step S77 from BEST and SECOND using the technique described hereinafter with reference to Figures 14 and 15. Each gene in this new member is mutated in step S78 based on the defined percentage chance of mutation using the uniformly distributed allele replacement. A new member is then evaluated in step S79 and in step S80 the new member is used to replace the WORST. In step S81 it is then determined whether the number of additional evaluations has been reached i.e. $a = A$ and if not the process returns to step S74. If the number of additional evaluations has been reached, the process proceeds to step S82 where all the members of the population are sorted based on the evaluation result with the lowest result as BEST. In step S83 the member of the population of the lowest evaluation result is then output. Details on this output member of the population can then be used for the configuration of the distributed database in order to improve the system performance.

Although in both genetic algorithms described above 5000 evaluations are used, any suitable number can be used. Mutation rate and population size can be appropriately selected to tune the genetic algorithm.

For example the mutation rate of 14% can be chosen and the population size of anything from 5 to 500.

The method of generating the new member in Figures 12 and 13 will now be described with reference to Figures 14 and 15.

Using the two parents, in step S90 an initial child is generated as an exact copy of parent 2. A portion of parent 1 of random length and at a random position is then selected in step S91 i.e. length = 5 and position = 8 in this example. This overlay portion is then overlaid onto a portion of the initial child of the same length at another random position in step S92 (i.e. at position = 4 in this example) to generate the resulting child as illustrated in Figure 15.

This technique is a variant of a two-point crossover technique which causes skewing. In this technique allele values in the child are directly overwritten by the overlay portion. There is no splicing and shunting of the genes.

This techniques is not only directly applicable to allelic representations in which the allowed allele range is the same for each gene, where allelic representations are not of the same range, mechanisms can be applied to align the allelic representations.

Figure 16 is a schematic illustration of a second distributed database scenario (B) wherein the nodes 400 and 500 are effectively split into two geographic regions

having low communications costs between nodes in the same region, and high costs between regions. As can be seen in the left hand side in Figure 16 the four nodes 400 comprising data servers 400a (A to D) and clients 400b (1 to 4) are interconnected via a communications network 600 of high speed. Similarly, in the right hand side the six nodes 500 comprising data servers 500a (F to J) and clients 500b (5 to 10) are interconnected via a communications network 700 of high speed. The two high speed communication networks 600 and 700 are interconnected via a low speed communications line 300. Each of the local communications network have a "supernode" whose data performance is 10 times that of the other nodes in the region.

Tables 4 to 6 below illustrate the modelling parameters used in the genetic algorithm methods for this scenario. A server contention overlap of 10% is assumed as for the first scenario (A).

Table 4

Base Comms Table (time in msec)
for Scenario B

Server /client	1	2	3	4	5	6	7	8	9	10
1	400	800	800	800	3000	3000	3000	3000	3000	3000
2	800	400	800	800	3000	3000	3000	3000	3000	3000
3	800	800	400	800	3000	3000	3000	3000	3000	3000
4	800	800	800	400	3000	3000	3000	3000	3000	3000
5	3000	3000	3000	3000	400	1000	1000	1000	1000	1000
6	3000	3000	3000	3000	1000	400	1000	1000	1000	1000
7	3000	3000	3000	3000	1000	1000	400	1000	1000	1000
8	3000	3000	3000	3000	1000	1000	1000	400	1000	1000
9	3000	3000	3000	3000	1000	1000	1000	1000	400	1000
10	3000	3000	3000	3000	1000	1000	1000	1000	1000	400

Table 5

Client application rates per node: Retrieval Rate (RR)
Update Rate (UR) Overlap (%) (in events per second)
for database Scenario B

Client	RR	UR	Overlap (%)
1	> 0.2	0.005	20
2	> 0.2	0.005	20
3	> 0.05	0.02	5
4	> 0.25	0.005	20
5	> 0.2	0.005	20
6	> 0.1	0.005	20
7	> 0.1	0.005	20
8	> 0.25	0.005	20
9	> 0.05	0.01	5
10	> 0.1	0.005	20

Table 6

Server Base Translation Time (BTT) per
Node (in msec) for database scenario B

Server	BTT
1	5000
2	5000
3	500
4	5000
5	5000
6	5000
7	5000
8	5000
9	500
10	5000

Figures 17 to 22 illustrate methods of determining an alternative cost value to the worst performance the cost value as determined in the method described with reference to Figure 11.

5 The method of Figure 17 is termed "just used" wherein the worst performance is output only for the servers which appear in the solution vector generated by the genetic algorithm.

10 In step S110 the Effective Transaction Rate (ETR) is calculated for each client by combining the retrieval and update rates, using the percentage overlap and equation 2 given hereinabove. In step S111 a two dimensional array L is populated with the update rate calculated for that client and indexed by client and
 15 server number. Where a server is to be used by a particular client as indicated in the input current proposed solution from the genetic algorithm (S113), in

step S112 with corresponding array entry in L is replaced with that clients Effective Transaction Rate. In step S114 the columns in the L array for servers which do not appear in the solution vector are then zeroed to remove their effect. A two dimensional array C is then created in step S115 to hold effective communication overhead delays. For each client/server interaction rate, now defined in L the corresponding entry in C is calculated using equation 1 given hereinabove. In step S116 the columns in the C array for servers which do not appear in the solution vector are then zeroed to remove their effect. In step S117 for each server the client entries in L are used to calculate aggregate loading for all clients on that server using equation 2 given hereinabove. In step S118 this is then converted to an average response time using equation 1 given hereinabove and this is stored in a one dimensional array T indexed by server. Any calculated infinite response times are replaced with a suitably large response time. In step S119, for each server which is represented in the solution vector, the client with the worst communication overhead delays is found from C and this is added to the relevant entry in T. In step S120 the largest value in T is then output and this represents the worst performance of the servers which are to be used by the clients in this database configuration.

Figure 18 illustrates another embodiment termed "plus average" in which 10% of the access times for all nodes is added to the least worst server time. This strikes a balance between minimising the worst performance and aggregate server performance.

In step S130 the Effective Transaction Rate (ETR) is calculated for each client by comparing the retrieval and update rates, using the percentage overlap and equation given hereinabove. In step S131 a two dimensional array L is populated with the update rate calculated for that client and indexed by client and server number. Where a server is to be used by a particular client as indicated in the input current proposed solution from the genetic algorithm (S133), in step 132 the corresponding array entry in L is replaced with that clients Effective Transaction Rate. A two dimensional array C is then created in step S134 to hold effective communication overhead delays. For each client/server interaction rate, now defined in L, the corresponding entry in C is calculated using equation 1 given hereinabove. In step S135, for each server, the client entries L are used to calculate aggregate loading for all clients on that server using equation 2 given hereinabove. In step S136 this is then converted into an average response time using equation 1 given hereinabove and this is stored in a one dimensional array T indexed by server. Any calculated infinite response

times are replaced with a suitably large response time. In step S137, for each server, the client with the worst communications overhead delay is found from C and this is added to the relevant entry in T. In step S138 10% of the average of all the values in T are calculated and in step S139 the largest value in T plus the 10% of the average of all of the average values in T is output as the performance measure for this database configuration.

Figure 19 is a flow diagram of another method of calculating a cost value for the distributed database. This model considers applying updates only to those nodes that are currently being accessed as servers, and adds 10% of the communications access time seen by all clients on the worst server, divided by the number of servers used. This adds a bias based on aggregate user perception but with weighting in favour of over duplication of data. This method can provide enhanced resilience and is referred to as "plus used".

In step S140 the Effective Transaction Rate (ETR) is calculated for each client by combining the retrieval and update rates, using the percentage overlap and equation 2 given hereinabove. In step S141 a two dimensional array L is populated with the update rate calculated for that client and indexed by client and server number. Where a server is to be used by a particular client as indicated in the input current proposed solution from the genetic algorithm (S143), in

step S142 the corresponding array entry in L is replaced with that clients Effective Transaction Rate (ETR). In step S144 the columns in the L array for the servers which do not appear in the solution vector are then
5 zeroed to remove their effect. In step S145 a two dimensional array C is then created to hold effective communication overhead delays. For each client/server interaction rate, now defined in L a corresponding entry in C is calculated using equation 1 given hereinabove.
10 In step S146 the columns in the C array for servers which do not appear in the solution vector is zeroed and in step S147 for each server which appears in the solution vector the client entries in L are used to calculate aggregate loading for all clients on that server using
15 equation 2 given hereinabove. In step S148 this is then converted to an average response time using equation 1 given hereinabove and this is stored in a one dimensional array T indexed by server. Any calculated infinite response times are replaced with a suitably large
20 response time. In step S152 for each server which is represented in the solution vector the client with the worst communications overhead delay is found from C and this is added to the relevant entry in T. Also in step S149 the worst server in T the average communication
25 times in C are calculated and in step S150 10% of the average is taken. This is then divided by the number of different servers appearing in the solution vector in

step S151 and in step S153 the largest value in T plus the value calculated in step S151 is output as the performance value for the database configuration.

Figure 20 is a flow diagram of another method of calculating a cost value for the distributed database. This method is a combination of "just used" and "plus average".

In step S160 the Effective Transaction Rate (ETR) is calculated for each client by combining the retrieval and update rates, using the percentage overlap and equation 2 given hereinabove. In step S161 a two dimensional array L is populated with the update rate calculated for that client and indexed by client and server number. Where a server is to be used by a particular client as indicated in the input current proposed solution from the genetic algorithm (S163), in step S162 the corresponding array entry in L is replaced by that clients Effective Transaction Rate (ETR). The columns in the L array for servers which do not appear in the solution vector are then zeroed in step S164 and in step S165 a two dimensional array C is then created to hold effective communication overhead delays. For each client/server interaction rate now defined in L the corresponding entry in C is calculated in equation 1 given hereinabove. The columns in the C array for servers which do not appear in the solution vector are then zeroed in step S166 and in step S167 for each used

server the client entries in L are used to calculate aggregate loading for all clients on that server using equation 2 given hereinabove. In step S168 this is then converted to an average response time using equation 1
5 given hereinabove and this is stored in a one dimensional array T indexed by server. Any calculated infinite response times are replaced with a suitably large response time. In step S169, for each server which is represented in the solution vector, the client with the
10 worst overhead communication delay is found from C and this is added to the relevant entry in T. In step S170 10% of the average of all of the values in T are calculated and in step S171 the largest value in T plus the 10% value calculated of the average of all of the
15 values in T is output as the performance value for the database.

Figure 21 is a flow diagram illustrating a further method of calculating a cost value. In this method only used servers i.e. servers which appear in the solution
20 vector the average of all client accesses weighted by their usage rate is added to all used servers. This technique is termed "plus all" and is thought to be the most realistic in terms of representing user perception of quality of service of the distributed database.

25 In step S180 the Effective Transaction Rate (ETR) is calculated for each client by combining the retrieval and update rates, using the percentage overlap and

equation 2 given hereinabove. In step S181 a two dimensional array L is populated with the update rate calculated for that client and indexed by the client and server number. Where a server is to be used by a particular client as indicated in the input current proposed solution from the genetic algorithm (S183), in step S182 the corresponding array entry in L is replaced by that clients Effective Transaction Rate (ETR). The columns in the L array for servers which do not appear in the solution vector are then zeroed in step S184. A two dimensional array C is then created in step S185 to hold effective communication overhead delays. For each client/server interaction rate, now defined in L the corresponding entry in C is calculated using equation 1 given hereinabove. The columns in the C array for servers which do not appear in the solution vector is then zeroed in step S186. The process then diverges to step S190 in which, for all used servers, the average of the communications times in C is calculated weighted by how often the communication links are used. Also in step S187 for each used server the client entries in L are used to calculate aggregate loading from all clients on that server using equation 2 given hereinabove. In step S188 this is then converted to an average response time using equation 1 given hereinabove and this is stored in a one dimensional array T indexed by server. Any calculated infinite response times are replaced with a

suitably large response time. In step S189 for each server which is represented in the solution vector the client with the worst communications overhead delay is found from C and this is added to the relevant entry in T. In step S191 the largest value in T plus the value calculated in step S190 are output as the performance value for the database.

Figure 22 is a flow diagram of a further method of calculating cost value. This method is similar to the "plus all" technique except that only 10% of the average of all clients accessed weighted by their usage rate is added. This technique is termed "plus 10%".

In step S200 the Effective Transaction Rate (ETR) is calculated for each client by combining the retrieval and update rates, using the percentage overlap and equation 2 given hereinabove. In step S201 a two dimensional array L is populated with the update rate calculated for that client and indexed by client and server number. Where a server is to be used by a particular client as indicated in the input current proposed solution from the genetic algorithm (S203), in step S202 the corresponding array entry in L is replaced with that clients Effective Transaction Rate. In step S204 the columns in the L array for servers which do not appear in the solution vector are zeroed. A two dimensional array C is then created in step S205 to hold effective communication overhead delays. For each

client/server interaction rate, now defined in L, the corresponding entry in C is calculated using equation 1 given hereinabove. In step S206 the columns in the C array for servers which do not appear in the solution vector are then zeroed and the process diverges. In step S210 for all used servers (i.e. the servers which appear in the solution vector) the average of the communications times in C weighted by how often the communication links are used are calculated. Also in step S207 for each server used, the client entries in L are used to calculate aggregate loading from all clients on that server using equation 2 given hereinabove. In step S208 this is then converted to an average response time using equation 1 given hereinabove and this is stored in a one dimensional array T indexed by server. Any calculated infinite response times are replaced by a suitably large response time. In step S209 for each server which is represented in the solution vector the client with the worst communications overhead delay is found from C and this added to the relevant entry in T. In step S211 the largest value in T plus 10% of the value calculated in step S210 is then output as the performance value for this database configuration.

So far the techniques described for obtaining the new solution vectors for use in the evaluation methods of Figures 11 and 17 to 25 have been evolutionary techniques wherein a genetic algorithm is used and thus

the technique of Figures 13 and 15 is used with two parent solution vectors. The present invention is not however limited to evolutionary techniques and is applicable to any search technique which can operate on a solution vector. Two non evolutionary techniques will now be described.

Figure 23 is a flow diagram of a Hill Climbing search technique wherein only one parent solution vector is used. In this non-evolutionary technique there is no "population".

In step S220 a random solution vector is generated initially and this becomes the current solution vector. In step S221 an iteration counter m is set to zero. In step S222 the current solution vector (a randomly generated solution vector) is evaluated using the methods of any one of Figures 11 and 17 to 22 to return a fitness value. The process then enters an iterative loop where in step S223 the iteration counter m is incremented. In step S224 it is determined whether the number of iterations have been complete i.e. $m = M$ where M is the total number of iterations to be carried out. If the number of iterations has been complete in step S225 the current solution vector is output as the optimum solution vector. If $m \neq M$, in step S226 a mutant solution vector is created from the current solution vector as will be described in more detail with reference to Figures 25 and 26. In step S227 the mutant solution vector is evaluated

to return a fitness value. The fitness value for the mutant solution vector is then compared with the fitness value for the current solution vector and if it is better then in step S229 the solution vector for the mutant is
5 set as the current solution vector and the process returns to step S223. If not the current solution vector is kept and the mutant solution vector is discarded and the process returns to step S223.

In accordance with this technique only improvements
10 to the current solution vector are kept. This technique, although simple, suffers from the disadvantage that if there is a localised minimum (or maximum if a maximum is to be found) in the search space, the Hill Climbing technique can determine the optimum solution to lie at
15 the localised minimum rather than at the global minimum.

Figure 24 is a flow diagram of another non-evolutionary search technique termed Simulated Annealing.

In step S230 a random solution vector is generated to become the current solution vector. In step S231 the
20 iteration counter m is set to 0. In step S232 the current solution vector is evaluated using the techniques of any one of Figures 11 and 17 to 22 in order to return a fitness value. In step S233 the process then enters an iteration loop wherein the iteration counter m is
25 incremented. In step S234 it is then determined whether the iteration loop is complete i.e. $m = M$, where M is the total number of iterations to be carried out. If so, in

step S235 the current solution vector is output as the optimum solution vector. If not in step S236 a mutant solution vector is created from the current solution vector as will be described in more detail hereinafter with references to Figures 25 and 26. In step S237 the mutant solution vector is then evaluated to return a fitness value. In step S238 it is determined whether the fitness value for the mutant solution vector is better than the fitness value for the current solution vector. If it is, in step S239, the solution vector for the mutant is set as the current solution and the process returns to step S233. If it is not, in step S240 the following calculation is made:

$$d = e^{\frac{(f_c - f_m)}{t}}$$

15

.....(5)

where f_c is fitness of current best solution,

f_m is fitness of the mutant,

and t is a "temperature" which starts high e.g. 10^6 and decays by a geometric cooling schedule.

20

In step S241 a random number n between 0 and 1 is then generated and in step S242 it is then determined whether $n > d$. If not in step S243 the solution vector for the mutant is set as the current solution vector and the process returns to step S233. If so the process returns

to step S233 retaining the current solution vector and discarding the mutant solution vector.

The Simulated Annealing technique has the advantage over the Hill Climbing technique in that a worst solution
5 can be accepted in the search procedure thereby allowing the search process to escape from localised minimum in the search for the global minimum.

The technique for generating the mutant solution vector will now be described with reference to Figures
10 25 and 26.

In step S301 for the single parent an initial child is generated as an exact copy of the parent. In step S302 from a random start position a portion of random length is selected from the parent as an overlay portion.
15 In Figure 26 it can be seen that the initial start position is 8 and the length of the overlay portion is 5. In step S303 the portion is then overlaid in the initial child on a portion of the initial child of the same length at a random start position to generate an
20 intermediate child. As can be seen in Figure 26 the random overlay position is 4. In step S304 a value in the child is then randomly mutated to generate a mutant solution vector to generate a resulting child. In Figure 26 the first value in the string is changed from A to G.
25

The use of the technique of the present invention in conjunction with a Tournament genetic algorithm has

been evaluated in comparison with a conventional Hill Climbing technique, a Simulated Annealing technique, a Breeder Genetic Algorithm technique using conventional uniform cross over to generate a new solution vector and
5 a Tournament genetic algorithm also using the uniform cross over technique to generate the new solution vector. Results have been obtained for scenario A as illustrated in Figure 9 and scenario B as illustrated in Figure 16.

Figure 27 shows the results for the five optimisers
10 for scenario A with the basic "least worst" server model. For each of 1000 runs, the fitness of the best solution was noted. The diagram firstly shows the percentage of these 1000 runs that found the known globally optimum solution value (referred to as "On Target" results). The
15 second group of columns shows the percentage of times that the fitness of the best found solution was within 5% of the known globally optimal solution (acceptable in industrial context). The third set of columns shows the percentage of times that the fitness of the best found
20 solution was more than 30% greater (given that lowest fitness value here is best) than the known globally optimal solution (deemed unacceptable in an industrial context). In Figure 27 to 29, the following abbreviations have been used:

25 HC - Hill Climbing technique
 SA - Simulated Annealing technique

- BDR - Breeder genetic algorithm technique using uniform cross over
- TNT - Tournament genetic algorithm technique using uniform cross-over
- 5 SKT - Tournament genetic algorithm using the technique of the present invention for generating a new solution vector.

As can be seen in Figure 27 whilst the tournament genetic algorithm using the technique of the present invention is well behind at 1000 evaluations, it still gives good results at 5000 evaluations.

Figure 28 shows the results for scenario A using the "plus all" model variant (seen as more indicative of user perceived quality of service). Once again, although the tournament genetic algorithm technique employing the technique of the present invention to generate a new solution vector gives unimpressive results at 1000 evaluations, good results were obtained at 5000 evaluations.

Figure 29 shows results for scenario B with the basic, least worst server model. Once again although the tournament genetic algorithm employing the technique of the present invention to generate the new solution vector is unimpressive at one thousand evaluations, the results are good for five thousand evaluations.

Figure 30 shows the results for scenario B with the "plus all" model. This problem clearly gives most

optimisers a large amount of difficulty. At 1000 evaluations, performance is wholly acceptable. However the tournament genetic algorithm technique utilising the technique of the present invention to generate the new solution vector provides goods results 5000 evaluations. Thus as can be seen in Figures 27 to 30, this technique is the only one which is able to give any degree of consistent performance which is critical in an industrial context.

10

Although the present invention has been described hereinabove with reference to its application to the optimization of the configuration of a distributed database, the technique of the present invention is widely applicable to the optimization of any physical system. The technique for generating a new solution vector exploits a particular feature of the chosen representation namely that good contiguous chunks of allele values in one part of the chromosome may also work well as a building block elsewhere in the chromosome.

15
20

The present invention can be applied to the optimisation of a distributed processing system of which a distributed data base is one example. In such a system applications (clients) requiring processing to be carried out at nodes (servers) in a network refer their processing requests to particular nodes. It is this configuration which can be adapted using the present

25

invention in order to distribute the processing load to improve the system performance. In the system, the operating parameters can comprise transaction or processing rates and network communication times can be used in the model as for the distributed data base embodiment.

The technique has also been applied to a hard benchmark problem in structural chemistry involving finding the two dimensional structure which minimizes the energy (as estimated by Lennard-Jones potentials) of a bonded string of atoms. Here the chromosome represents a list of adjacent bond angles and the true optimum in each case is a close-to spiral structure in which the adjacent angles are similar to each other (but not exactly the same), with repeated contiguous patterns of angles along the string. Coded as a maximisation problem, tables 7 and 8 below indicate the results over ten trials for both ten and twenty atom cases for a tournament genetic algorithm using one point (1 pt) two point (2 pt), uniform (unif), and the inventive cross-over technique (skew).

Table 7

	10 atoms:									
1 pt:	20	19	19	19	19	19	19	19	19	19
2 pt:	19	18	19	18	20	19	18	41	19	19
unif:	19	20	19	17	19	18	18	19	19	20
skew:	20	20	20	20	20	20	20	20	20	20

Table 8

	20 atoms:									
1 pt:	37	40	38	37	41	41	20	42	37	39
2 pt:	41	39	40	39	40	41	40	40	41	42
unif:	43	43	37	39	41	41	39	42	41	42
skew:	47	47	45	47	47	46	46	45	47	47

It can be seen from tables 7 and 8 that the technique of the present invention provides the best solution to this problem both in terms of best value found and repeatability of results indicating the wider applicability of this technique.

A further application of this technique is to the configuration of switched networks wherein information units comprise call routing information for controlling

the switching of network switches to route calls. An example of where usage would change in such a network is where a company decides to offer a special discount on calls in an under utilised part of the communications network at a time later that day. A system could adapt itself to allow efficient access of the appropriate data for the duration of the discount in advance and show that the system is optimised to meet the expected demand. In another example, in a mobile telephone network where a customer moves from one city to another, their data records can follow so that they are close by rather than being accessed slowly across large distances. This can take place by monitoring retrieval performance of the data for the customer and adapting the distribution of data to include the retrieval performance.

Clearly the present invention has a wide range of applications and is particularly suited to optimization techniques with complex cost-values which have a complex search space e.g. more than one minimum or maximum.

The application of the technique in genetic algorithms can be tailored to suit the problem by selecting the mutation rate and population size accordingly.

As has been shown the present invention is particularly suited to optimizing a quality of service metric in a distributed database taking into account both

load on the servers and loads on the links interconnecting the servers and clients.

Although the present invention has been described hereinabove with reference to allelic representations of the solution vector, the present invention is not limited to such representations and is applicable to canonical representations.

The present invention can be implemented as a computer program operating on a standard computer and thus the present invention can be embodied as a storage medium containing processor implementable instructions for controlling a processor to carry out the method. Further, since the computer code can be transmitted in electronic form for example by being downloaded over a network, the present invention can be embodied as an electronic signal carrying computer code for controlling a computer to carry out the method.

Although the present invention has been described hereinabove with reference to specific embodiments, the present invention is not limited to such embodiments and it would be apparent to a skilled person in the art that modifications can be made within the spirit and scope of the appended claims.

CLAIMS:

1. Data processing apparatus for determining optimum parameters of a model of a physical system, the apparatus
5 comprising:

obtaining means for obtaining at least one initial string of values representing the parameters of the model to be optimised;

10 first determining means for determining a cost value associated with the model having parameters represented by the or each string of values;

generating means for repeatedly generating a new string of values by selecting a sequence of values of random length, starting at a random position in a said
15 string of values, replacing a sequence of values of the same length in a said string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values;

20 wherein said determining means is adapted to determine a cost value associated with the model having parameters represented by the new string of values; and

second determining means for determining the optimum parameters as one of said initial or new string of values
25 for which the cost value associated with the model having the optimum parameters is closest to a target.

61

2. Data processing apparatus according to claim 1 wherein said generating means is adapted to consider a last value in a said string of values as being sequential to a first value in said string of values such that said
5 selected sequence of values can include said last and first values sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

10

3. Data processing apparatus according to claim 1 or claim 2, wherein said obtaining means is adapted to randomly obtain values for said at least one initial string of values.

15

4. Data processing apparatus according to any preceding claim wherein said obtaining means is adapted to obtain a plurality of said initial strings of values to form a population, and said generating means is adapted to
20 perform a genetic algorithm on said population by repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in a first said string of values,

25

(b) replacing a sequence of values of the same length in a second said string of values at a random position, and

(c) mutating the value of one or more of the values of the resulting string of values to generate said new string of values.

5 5. Data processing apparatus according to claim 4, wherein said generating means is adapted to repeatedly

(a) delete a proportion of the population for which the cost values are furthest from said target, and

(b) perform said repeated generation of new strings
10 of values until the population is regenerated; and

to select said first and second strings of values randomly from the undeleted proportion of the population for which the cost values are closest to said target.

15

6. Data processing apparatus according to claim 4, wherein said generating means is adapted to repeatedly

(a) randomly select three said strings of values from said population,

20 (b) select two of the selected strings of values for which the cost values are closest to the target as said first and second strings of values, and

(c) replace the third of said selected strings of values with the generated new string of values.

25

7. Data processing apparatus according to any one of claims 1 to 3, wherein said obtaining means is adapted

to obtain a single said initial string of values as a parent string of values and said generating means is adapted to repeatedly generating a new string of values by

5 (a) selecting a sequence of values of random length starting at a random position in said parent string of values,

 (b) replacing a sequence of values of the same length in said parent string of values at a random
10 position, and

 (c) changing the value of one or more of the values of the resulting string of values to generate said new string of values.

15 8. Data processing apparatus according to claim 7, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values if the cost value associated with the model having said new string of parameters is closer to said
20 target.

 9. Data processing apparatus according to claim 8 wherein said generating means is adapted to repeatedly replace said parent string of values with said new string
25 of values if the exponential of the difference in the cost values for said parent string of values and said new string of values divided by a factor dependent upon the

number of repetitions by said generating means is greater than a random number between 0 and 1.

10. Apparatus for controlling the configuration of a physical system, the apparatus comprising:

obtaining means for obtaining at least one initial string of values representing configuration parameters of the physical system;

modelling means for modelling the physical system using said configuration parameters to determine a cost value for the or each initial string of values;

generating means for repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values; wherein said modelling means is adapted to determine a cost value for the configuration parameters represented by said new string of values;

determining means for determining optimum configuration parameters represented as one of said initial or new string of values for which the cost value is closest to a target; and

control means for controlling the configuration of said physical system in accordance with the determined optimum configuration parameters.

5 11. Apparatus according to claim 10, wherein said generating means is adapted to consider a last value in a said string of values as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values
10 sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

12. Apparatus according to claim 10 or claim 11, wherein
15 said obtaining means is adapted to randomly obtain values for said at least one initial string of values.

13. Apparatus according to any one of claims 10 to 12, wherein said obtaining means is adapted to obtain a
20 plurality of said initial strings of values to form a population, and said generating means is adapted to perform a genetic algorithm on said population by repeatedly generating a new string of values by

(a) selecting a sequence of values of random length
25 starting at a random position in a first said string of values,

(b) replacing a sequence of values of the same length in a second said string of values at a random position, and

(c) mutating the value of one or more of the values
5 of the resulting string of values to generate said new string of values.

14. Apparatus according to claim 13, wherein said generating means is adapted to repeatedly

10 (a) delete a proportion of the population for which the cost values are furthest from said target, and

(b) perform said repeated generation of new strings of values until the population is regenerated; and

15 to select said first and second strings of values randomly from the undeleted proportion of the population for which the cost values are closest to said target.

15. Apparatus according to claim 13, wherein said
20 generating means is adapted to repeatedly

(a) randomly select three said strings of values from said population,

(b) select two of the selected strings of values for which the cost values are closest to the target as
25 said first and second strings of values and

(c) replace the third of said selected strings of values with the generated new string of values.

16. Apparatus according to any one of claims 10 to 12, wherein said obtaining means is adapted to obtain a single said initial string of values as a parent string of values and said generating means is adapted to repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in said parent string of values,

(b) replacing a sequence of values of the same length in said parent string of values at a random position, and

(c) changing the value of one or more of the values of the resulting string of values to generate said new string of values.

17. Apparatus according to claim 16, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values if the cost value associated with the model having said new string of parameters is closer to said target.

18. Apparatus according to claim 16, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values if the exponential of the difference in the cost values for said parent string of values and said new string of values divided by a factor dependent upon the number of

repetitions by said generating means is greater than a random number between 0 and 1.

19. Apparatus according to any one of claims 10 to 18 including means for receiving operating parameters from said physical system; wherein said model means is adapted to model said physical system using the received operating parameters.

20. A distributed database controller for configuring a distributed database comprising a plurality of data servers connected over a network, each data server holding any number of data units, and a plurality of clients connected over the network to the data servers, each of said clients being adapted to retrieve data units from and/or update data units at one or more of the data servers, the controller comprising:

means for obtaining at least one initial string of values representing configuration parameters for the distributed database defining which data server is to be accessed by which client;

modelling means for modelling the distributed database using the configuration parameters and information on the passage of data within the distributed database to determine a performance value for the or each initial string of values;

generating means for repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position and changing the value of one or more of the values of the resulting string of values to generate a new string of values, wherein said modelling means is adapted to determine a performance value for the configuration parameters represented by said new string of values;

determining means for determining optimum configuration parameters represented as one of said initial or new string of values for which the performance value is the best; and

control means for configuring the distributed database to control which data server is to be accessed by which client in accordance with the determined optimum configuration parameters.

20

21. A controller according to claim 20, wherein said generating means is adapted to consider a last value in a said string of values as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values sequentially in a said string of values and the sequence

25

of values to be replaced can include said last and first values sequentially in a said string of values.

22. A controller according to claim 20 or claim 21,
5 wherein said obtaining means is adapted to randomly obtain values for said at least one initial string of values.

23. A controller according to any one of claims 20 to
10 22, wherein said obtaining means is adapted to obtain a plurality of said initial strings of values to form a population, and said generating means is adapted to perform a genetic algorithm on said population by repeatedly generating a new string of values by

15 (a) selecting a sequence of values of random length starting at a random position in a first said string of values,

(b) replacing a sequence of values of the same length in a second said string of values at a random
20 position, and

(c) mutating the value of one or more of the values of the resulting string of values to generate said new string of values.

25 24. A controller according to claim 23, wherein said generating means is adapted to repeatedly

(a) delete a proportion of the population for which the performance values are worst, and

(b) perform said repeated generation of new strings of values until the population is regenerated; and

5 to select said first and second strings of values randomly from the undeleted proportion of the population for which the performance values are best.

25. A controller according to claim 23, wherein said
10 generating means is adapted to repeatedly

(a) randomly select three said strings of values from said population,

(b) select two of the selected strings of values for which the performance values are worst, and

15 (c) replace the third of said selected strings of values with the generated new string of values.

26. A controller according to any one of claims 20 to 22, wherein said obtaining means is adapted to obtain a
20 single said initial string of values as a parent string of values and said generating means is adapted to repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in said parent string of
25 values,

(b) replacing a sequence of values of the same length in said parent string of values at a random position, and

(c) changing the value of one or more of the values
5 of the resulting string of values to generate said new string of values.

27. A controller according to claim 26, wherein said
10 generating means is adapted to repeatedly replace said parent string of values with said new string of values if the performance value associated with the model having said new string of parameters is best.

28. A controller according to claim 26, wherein said
15 generating means is adapted to repeatedly replace said parent string of values with said new string of values if the exponential of the difference in the performance values for said parent string of values and said new string of values divided by a factor dependent upon the
20 number of repetitions by said generating means is greater than a random number between 0 and 1.

29. A controller according to any one of claims 20 to
25 28, wherein said control means is adapted to copy, move and/or delete data units at data servers to change the distribution of the data units across the data servers and/or to change the data servers to be accessed by said

clients using the determined optimum configuration parameters to improve the system performance.

30. A controller according to any one of claims 20 to
5 29, wherein said modelling means is adapted to determine said performance value in dependence upon the time taken for one or more said clients to retrieve and/or update data units at one or more data servers.

10 31. A controller according to claim 30 wherein said modelling means is adapted to determine said performance value as the longest response time for any said client to access any said data server.

15 32. A controller according to claim 30, wherein said modelling means is adapted to determine said performance value as the longest response time for any said client to access any said data server and a proportion of the average of the response times for all said data servers.

20
33. A controller according to claim 30, wherein said modelling means is adapted to determine said performance value as a function of the rates at which data can be retrieved by said clients from said data servers, the
25 rates at which data can be updated by said clients at said data servers, the contention between said clients

for accessing said data servers, and communications times between said clients and said data servers.

34. A controller according to any one of claims 20 to
5 33 including monitoring means for monitoring the passage of data over the network, wherein said modelling means is adapted to use the output of said monitoring means to determine said performance values, and said control means is adapted to adaptively configure the distributed
10 database.

35. A distributed processing controller for configuring a distributed processing system comprising a plurality of servers connected over a network, each server being
15 capable of carrying out any number of processing operations, and a plurality of clients connected over the network to the servers, each of said clients being adapted to request processing to be carried out by one or more of the servers, the controller comprising:

20 means for obtaining at least one initial string of values representing configuration parameters for the distributed processing system defining which server is to be accessed by which client;

25 modelling means for modelling the distributed processing system using the configuration parameters and information on the communications speed within the

distributed processing system to determine a performance value for the or each initial string of values;

generating means for repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position and changing the value of one or more of the values of the resulting string of values to generate a new string of values, wherein said modelling means is adapted to determine a performance value for the configuration parameters represented by said new string of values;

determining means for determining optimum configuration parameters represented as one of said initial or new string of values for which the performance value is the best; and

control means for configuring the distributed processing system to control which server is to be accessed by which client in accordance with the determined optimum configuration parameters.

36. A controller according to claim 35, wherein said generating means is adapted to consider a last value in a said string of values as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values

sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

5 37. A controller according to claim 35 or claim 36, wherein said obtaining means is adapted to randomly obtain values for said at least one initial string of values.

10 38. A controller according to any one of claims 35 to 37, wherein said obtaining means is adapted to obtain a plurality of said initial strings of values to form a population, and said generating means is adapted to perform a genetic algorithm on said population by
15 repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in a first said string of values,

(b) replacing a sequence of values of the same
20 length in a second said string of values at a random position, and

(c) mutating the value of one or more of the values of the resulting string of values to generate said new string of values.

25

39. A controller according to claim 38, wherein said generating means is adapted to repeatedly

(a) delete a proportion of the population for which the performance values are worst, and

(b) perform said repeated generation of new strings of values until the population is regenerated; and

5 to select said first and second strings of values randomly from the undeleted proportion of the population for which the performance values are best.

40. A controller according to claim 38, wherein said
10 generating means is adapted to repeatedly

(a) randomly select three said strings of values from said population,

(b) select two of the selected strings of values for which the performance values are worst, and

15 (c) replace the third of said selected strings of values with the generated new string of values.

41. A controller according to any one of claims 35 to 37, wherein said obtaining means is adapted to obtain a
20 single said initial string of values as a parent string of values and said generating means is adapted to repeatedly generating a new string of values by

(a) selecting a sequence of values of random length starting at a random position in said parent string of
25 values,

(b) replacing a sequence of values of the same length in said parent string of values at a random position, and

(c) changing the value of one or more of the values
5 of the resulting string of values to generate said new string of values.

42. A controller according to claim 41, wherein said
10 generating means is adapted to repeatedly replace said parent string of values with said new string of values if the performance value associated with the model having said new string of parameters is best.

43. A controller according to claim 41, wherein said
15 generating means is adapted to repeatedly replace said parent string of values with said new string of values if the exponential of the difference in the performance values for said parent string of values and said new string of values divided by a factor dependent upon the
20 number of repetitions by said generating means is greater than a random number between 0 and 1.

44. A controller according to any one of claims 35 to
25 43, wherein said modelling means is adapted to determine said performance value in dependence upon the transaction times by said servers.

45. A controller according to claim 44 wherein said modelling means is adapted to determine said performance value as the longest transaction time by any said server.

5 46. A controller according to claim 44, wherein said modelling means is adapted to determine said performance value as the longest transaction time by any said data server and a proportion of the average of the transaction times for all said servers.

10

47. A controller according to any one of claims 35 to 46 including monitoring means for monitoring the transaction times of said servers, wherein said modelling means is adapted to use the output of said monitoring means to determine said performance values, and said control means is adapted to adaptively configure the distributed processing system.

15 48. A processor implemented method of determining optimum parameters of a model of a physical system, the method comprising:

obtaining at least one initial string of values representing the parameters of the model to be optimised;

25 determining a cost value associated with the model having parameters represented by the or each string of values;

repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a
5 said string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values;

determining a cost value associated with the model having parameters represented by the new string of
10 values; and

determining the optimum parameters as one of said initial or new string of values for which the cost value is closest to a target.

15 49. A method according to claim 48, wherein in the generating step a last value in a said string of values is considered as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values
20 sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

50. A method according to claim 48 or claim 49, wherein
25 the values for said at least one initial string of values is randomly obtained.

81

51. A method according to any one of claims 48 to 50, wherein a plurality of said initial strings of values are obtained to form a population, and the repeated generating step comprises selecting a sequence of values
5 of random length starting at a random position in a first said string of values, replacing a sequence of values of the same length in a second said string of values, and changing the value of one or more of the values of the resulting string of values to generate said new string
10 of values.

52. A method according to claim 51, including repeatedly:

(a) deleting a proportion of the population for
15 which the cost value is furthest from said target;

(b) performing said repeated generating step until the population is regenerated; and

(c) selecting said first and second strings of values randomly from the proportion of the population for
20 which the cost values are closest to said target.

53. A method according to claim 51 including repeatedly:

(a) randomly selecting three said strings of values from said population

25 (b) selecting two of the selected strings of values for which the cost values are closest to said target as

said first and second strings of values for use in the generating step; and

(c) replacing the third of said selected string of values with the generated new string of values.

5

54. A method according to any one of claims 48 to 50, wherein a single said initial string of values is obtained as a parent string of values, and the repeated generating step comprises selecting a sequence of values or random length starting at a random position in said parent string of values, replacing a sequence of values of the same length in said parent string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate said new string of values.

10
15

55. A method according to claim 54, including repeatedly replacing said parent string of values with said new string of values if the cost value for said new string of values is closer to said target.

20

56. A method according to claim 55, including repeatedly replacing said parent string of values with said new string of values if the exponential of the difference in the cost values for said parent and new strings of values divided by as factor dependent upon the number of

25

repetitions of the generating step is greater than a randomly generated number between 0 and 1.

57. A method of controlling the configuration of a physical system, the method comprising:

obtaining at least one initial string of values representing configuration parameters of the physical system;

10 modelling the physical system using said configuration parameters to determine a cost value for the or each initial string of values;

15 repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position and changing the value of one or more of the values of the resulting string of values to generate a new string of values;

20 modelling the physical system using said configuration parameters to determine cost values for the new strings of values;

determining optimum configuration parameters represented as one of said initial or new string of values for which the cost value is closest to a target;
25 and

controlling the configuration of said physical system in accordance with the determined optimum configuration parameters.

5 58. A method according to claim 57, wherein in the generating step a last value in a said string of values is considered as being sequential to a first value in said string of values such that said selected sequence of values can include said last and first values
10 sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

59. A method according to claim 57 or claim 58, wherein
15 the values for said at least one initial string of values is randomly obtained.

60. A method according to any one of claims 57 to 59, wherein a plurality of said initial strings of values are
20 obtained to form a population, and the repeated generating step comprises selecting a sequence of values of random length starting at a random position in a first said string of values, replacing a sequence of values of the same length in a second said string of values, and
25 changing the value of one or more of the values of the resulting string of values to generate said new string of values.

61. A method according to claim 60, including repeatedly:

(a) deleting a proportion of the population for which the cost value is furthest from said target;

5 (b) performing said repeated generating step until the population is regenerated; and

(c) selecting said first and second strings of values randomly from the proportion of the population for which the cost values are closest to said target.

10

62. A method according to claim 61 including repeatedly:

(a) randomly selecting three said strings of values from said population;

15 (b) selecting two of the selected strings of values for which the cost values are closest to said target as said first and second strings of values for use in the generating step; and

(c) replacing the third of said selected string of values with the generated new string of values.

20

63. A method according to any one of claims 57 to 59, wherein a single said initial string of values is obtained as a parent string of values, and the repeated generating step comprises selecting a sequence of values
25 of random length starting at a random position in said parent string of values, replacing a sequence of values of the same length in said parent string of values at a

random position, and changing the value of one or more of the values of the resulting string of values to generate said new string of values.

- 5 64. A method according to claim 63, including repeatedly replacing said parent string of values with said new string of values if the cost value for said new string of values is closer to said target.
- 10 65. A method according to claim 63, including repeatedly replacing said parent string of values with said new string of values if the exponential of the difference in the cost values for said parent and new strings of values divided by as factor dependent upon the number of
15 repetitions of the generating step is greater than a randomly generated number between 0 and 1.
- 20 66. A method according to any one of claims 57 to 65, including receiving operating parameters from said physical system, wherein the modelling step includes using the received operating parameters to model said physical system.
- 25 67. A method of configuring a distributed database comprising a plurality of data servers connected over a network, each data server holding any number of data units, and a plurality of clients connected over said

network to said data servers, each of said client being adapted to retrieve data units from and/or update data units at one or more of the data servers, the method comprising:

5 obtaining at least one initial string of values representing configuration parameters for the distributed database defining which data server is to be accessed by which client;

10 modelling the distributed database using the configuration parameters and information on communications within the distributed database to determine a performance value for the or each initial string of values;

15 repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values, replacing a sequence of values of the same length in a said string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values,

20 modelling the distributed database using the configuration parameters and information on the passage of data within the distributed database to determine a performance value for the new strings of values;

25 determining optimum configuration parameters represented as one of said initial or new string of values for which the performance value is the best; and

configuring the distributed database to control which data server is to be accessed by which client in accordance with the determined optimum configuration parameters.

5

68. A method according to claim 67, wherein in the generating step a last value in a said string of values is considered as being segmented to a first value in said string of values such that said selected sequence of values can include said last and first values sequentially in a said string of values and the sequence of values to be replaced can include said last and first values sequentially in a said string of values.

10 69. A method according to claim 67 or claim 68, wherein the values for said at least one initial string of values is randomly obtained.

15 70. A method according to any one of claims 67 to 69, wherein a plurality of said initial strings of values are obtained to form a population, and the repeated generating step comprises selecting a sequence of values of random length starting at a random position in a first said string of values, replacing a sequence of values of the same length in a second said string of values, and changing the value of one or more of the values of the

20

25

resulting string of values to generate said new string of values.

71. A method according to claim 70, including
5 repeatedly:

(a) deleting a proportion of the population for which the performance values are worst;

(b) performing said repeated generating step until the population is regenerated; and

10 (c) selecting said first and second strings of values randomly from the proportion of the population for which the performance values are best.

72. A method according to claim 70, including
15 repeatedly:

(a) randomly selecting three said strings of values from said population;

(b) selecting two of the selected strings of values for which the performance values are best; and

20 (c) replacing the third of said selected string of values with the generated new string of values.

73. A method according to any one of claims 67 to 69, wherein a single said initial string of values is
25 obtained as a parent string of values, and the repeated generating step comprises selecting a sequence of values of random length starting at a random position in said

parent string of values, replacing a sequence of values of the same length in said parent string of values at a random position, and changing the value of one or more of the values of the resulting string of values to
5 generate said new string of values.

74. A method according to claim 73, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values
10 if the performance value associated with the model having said new string of parameters is best.

75. A method according to claim 73, wherein said generating means is adapted to repeatedly replace said
15 parent string of values with said new string of values if the exponential of the difference in the performance values for said parent string of values and said new string of values divided by a factor dependent upon the number of repetitions by said generating means is greater
20 than a random number between 0 and 1.

76. A method according to any one of claims 67 to 75, wherein the step of configuring the distributed database comprises copying, moving and/or deleting data units at
25 data servers to change the distribution of data units across the data servers, and/or changing the data servers to be accessed by said clients using the determined

optimum configuration parameters to improve the system performance.

5 77. A method according to any one of claims 67 to 76, wherein said performance values is determined in dependence upon the time taken for one or more said clients to retrieve and/or update data units at one or more data servers.

10 78. A method according to claim 77, wherein the performance value is determined as the longest response time for any said client means to access any said data server.

15 79. A method according to claim 77, wherein said performance value is determined as the longest response time for any said data server and a proportion of the average of the response times for all said data servers.

20 80. A method according to claim 77, wherein said performance value is determined as a function of the rates at which data can be retrieved by said clients from said data servers, the rates at which data can be updated by said clients at said data servers, the contention
25 between said clients for accessing said data servers, and communication times between said clients and said data servers.

81. A method according to any one of claims 67 to 80 including the step of monitoring the passage of information over the network, wherein the performance value is determined in accordance with the monitoring, and the configuration of the distributed database is adaptively controlled.

82. A method of configuring a distributed processing system comprising a plurality of servers connected over a network, each server being capable of carrying out any number of processing operations, and a plurality of clients connected over said network to said servers, each of said client being adapted to request processing to be carried out by one or more of the servers, the method comprising:

obtaining at least one initial string of values representing configuration parameters for the distributed processing system defining which server is to be accessed by which client;

modelling the distributed processing system using the configuration parameters and information on communications within the distributed database to determine a performance value for the or each initial string of values;

repeatedly generating a new string of values by selecting a sequence of values of random length starting at a random position in a said string of values,

replacing a sequence of values of the same length in a said string of values at a random position, and changing the value of one or more of the values of the resulting string of values to generate a new string of values,

5 modelling the distributed processing system using the configuration parameters and information on the passage of data within the distributed processing system to determine a performance value for the new strings of values;

10 determining optimum configuration parameters represented as one of said initial or new string of values for which the performance value is the best; and
 configuring the distributed processing system to control which server is to be accessed by which client
15 in accordance with the determined optimum configuration parameters.

83. A method according to claim 82, wherein in the generating step a last value in a said string of values
20 is considered as being segmented to a first value in said string of values such that said selected sequence of values can include said last and first values sequentially in a said string of values and the sequence of values to be replaced can include said last and first
25 values sequentially in a said string of values.

84. A method according to claim 82 or claim 83, wherein the values for said at least one initial string of values is randomly obtained.

5 85. A method according to any one of claims 82 to 84, wherein a plurality of said initial strings of values are obtained to form a population, and the repeated generating step comprises selecting a sequence of values of random length starting at a random position in a first
10 said string of values, replacing a sequence of values of the same length in a second said string of values, and changing the value of one or more of the values of the resulting string of values to generate said new string of values.

15

86. A method according to claim 85, including repeatedly:

(a) deleting a proportion of the population for which the performance values are worst;

20 (b) performing said repeated generating step until the population is regenerated; and

(c) selecting said first and second strings of values randomly from the proportion of the population for which the performance values are best.

25

87. A method according to claim 85, including repeatedly:

(a) randomly selecting three said strings of values from said population;

(b) selecting two of the selected strings of values for which the performance values are best; and

5 (c) replacing the third of said selected string of values with the generated new string of values.

10 88. A method according to any one of claims 82 to 84, wherein a single said initial string of values is obtained as a parent string of values, and the repeated generating step comprises selecting a sequence of values of random length starting at a random position in said parent string of values, replacing a sequence of values of the same length in said parent string of values at a
15 random position, and changing the value of one or more of the values of the resulting string of values to generate said new string of values.

20 89. A method according to claim 88, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values if the performance value associated with the model having said new string of parameters is best.

25 90. A method according to claim 88, wherein said generating means is adapted to repeatedly replace said parent string of values with said new string of values

if the exponential of the difference in the performance values for said parent string of values and said new string of values divided by a factor dependent upon the number of repetitions by said generating means is greater
5 than a random number between 0 and 1.

91. A method according to any one of claims 82 to 90, wherein said performance values is determined in dependence upon the transaction time taken for said
10 servers.

92. A method according to claim 91, wherein the performance value is determined as the longest transaction time for said data server.
15

93. A method according to claim 91, wherein said performance value is determined as the longest transaction time for any said server and a proportion of the average of the transaction times for all said
20 servers.

94. A method according to any one of claims 82 to 93 including the step of monitoring the transaction times of said servers, wherein the performance value is
25 determined in accordance with the monitoring, and the configuration of the distributed processing system is adaptively controlled.

97

95. A storage medium storing processor implementable instructions for controlling a processor to carry out the method of any one of claims 48 to 94.
- 5 96. An electronic signal carrying computer code for controlling a processor to carry out the method of any one of claims 48 to 94.




Fig 1

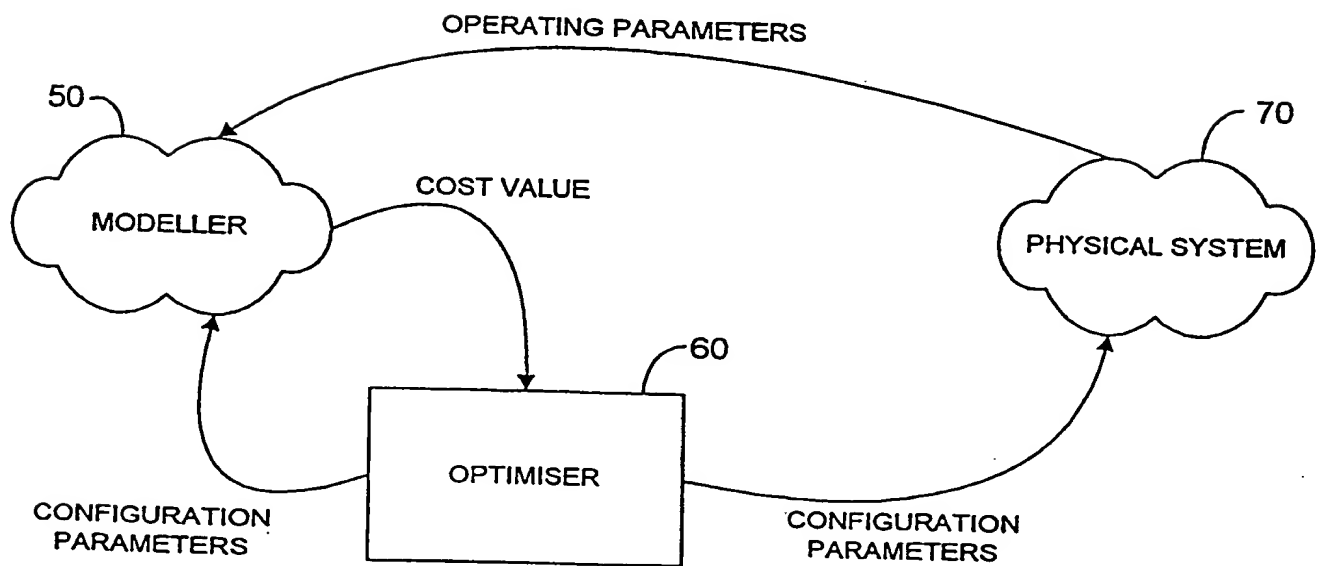


Fig 2

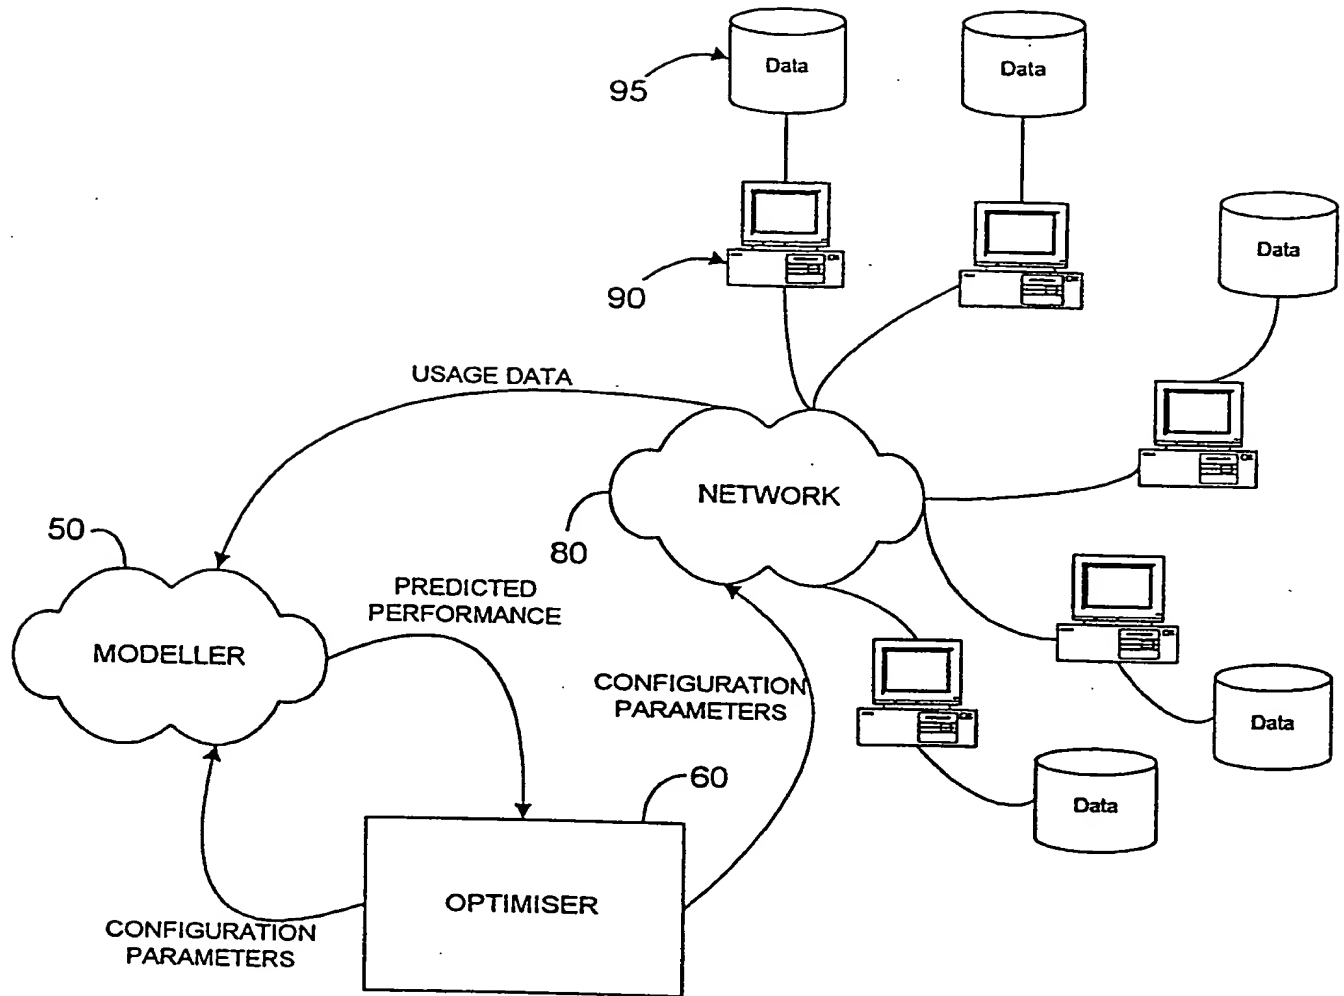


Fig 3

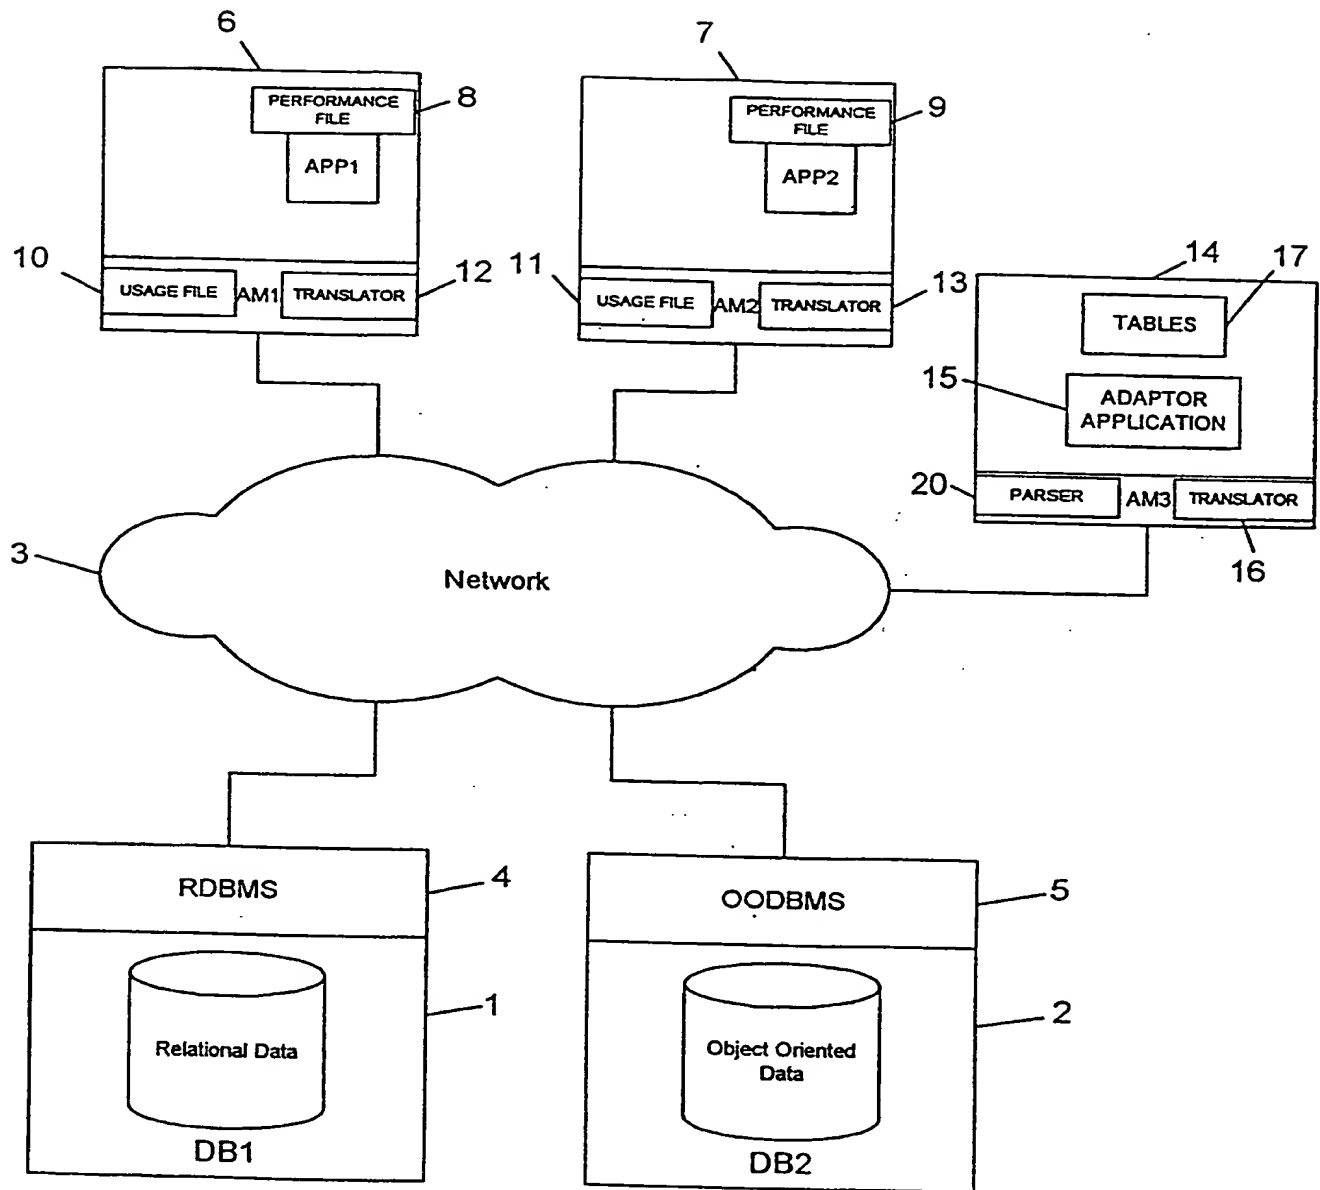


Fig 4

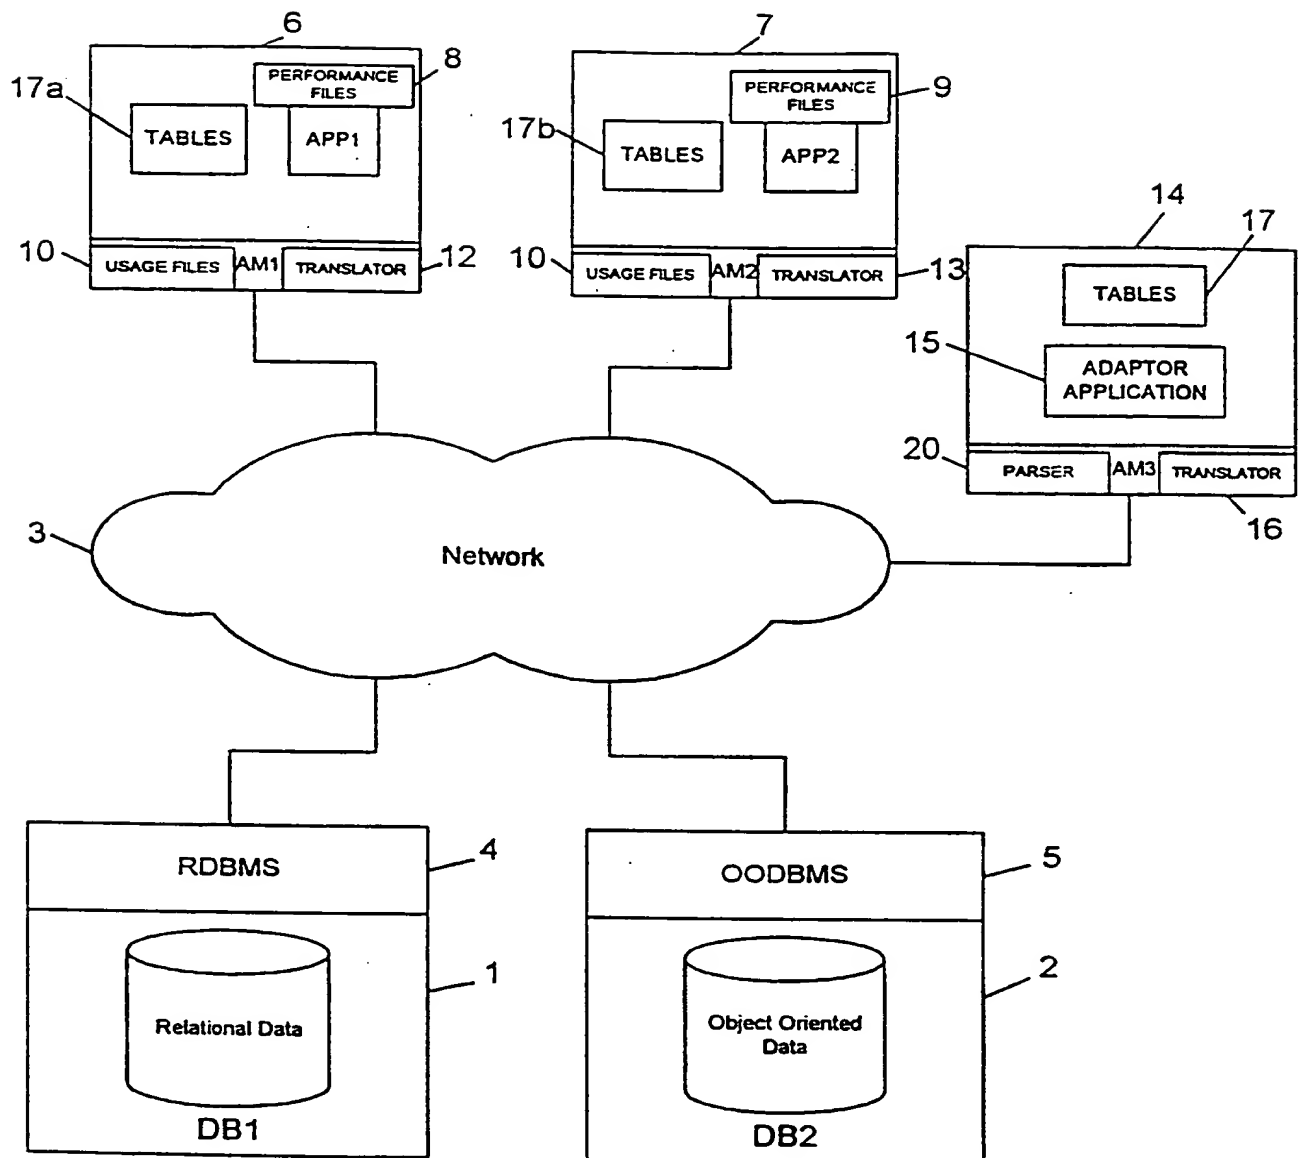


Fig 5




Fig 6

Fig 7a

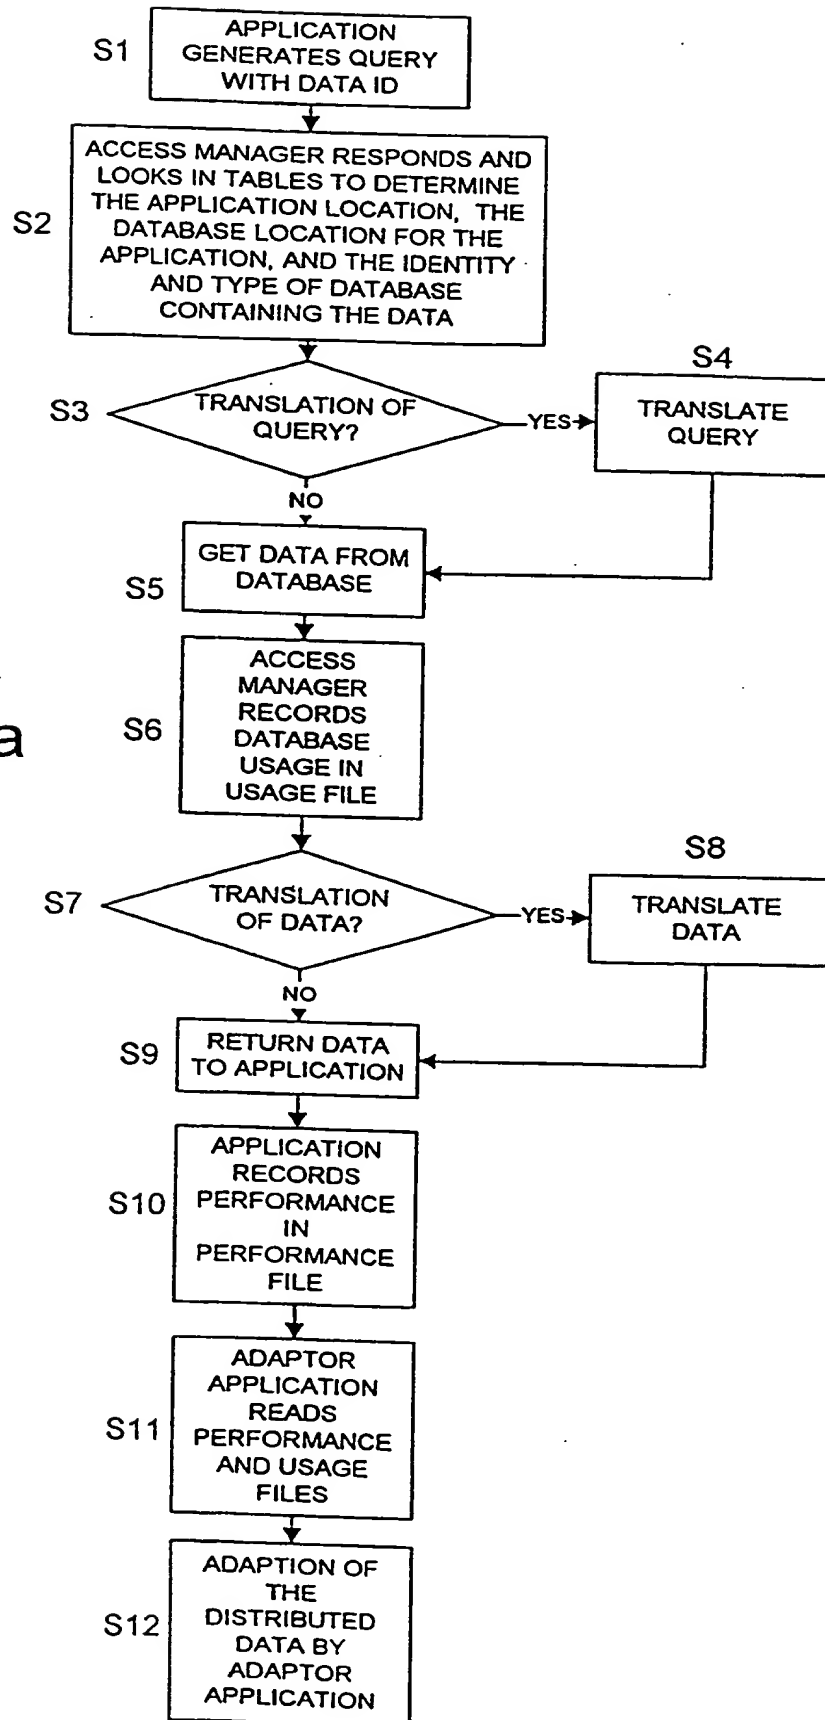


Fig 7b

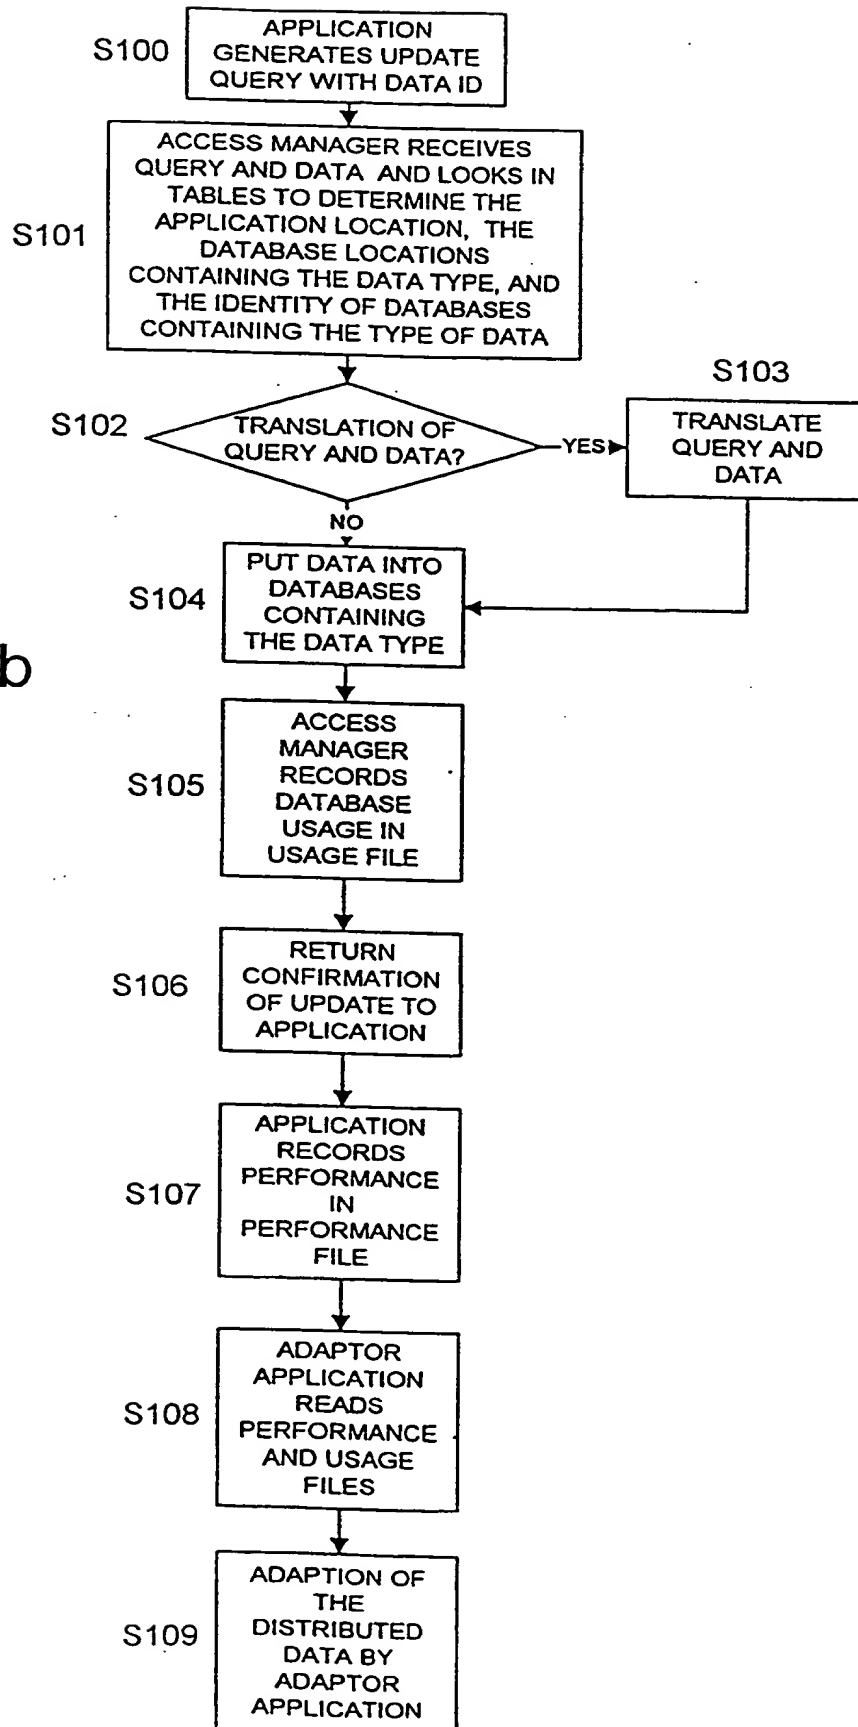
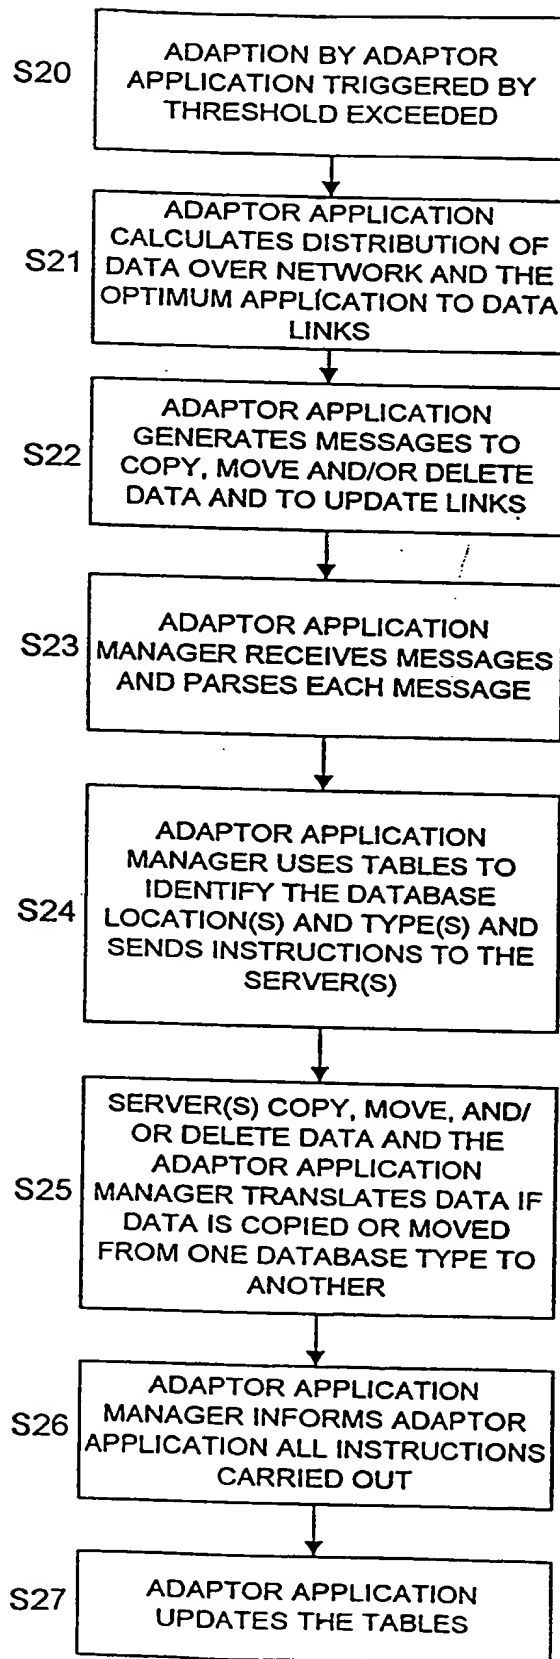


Fig 8




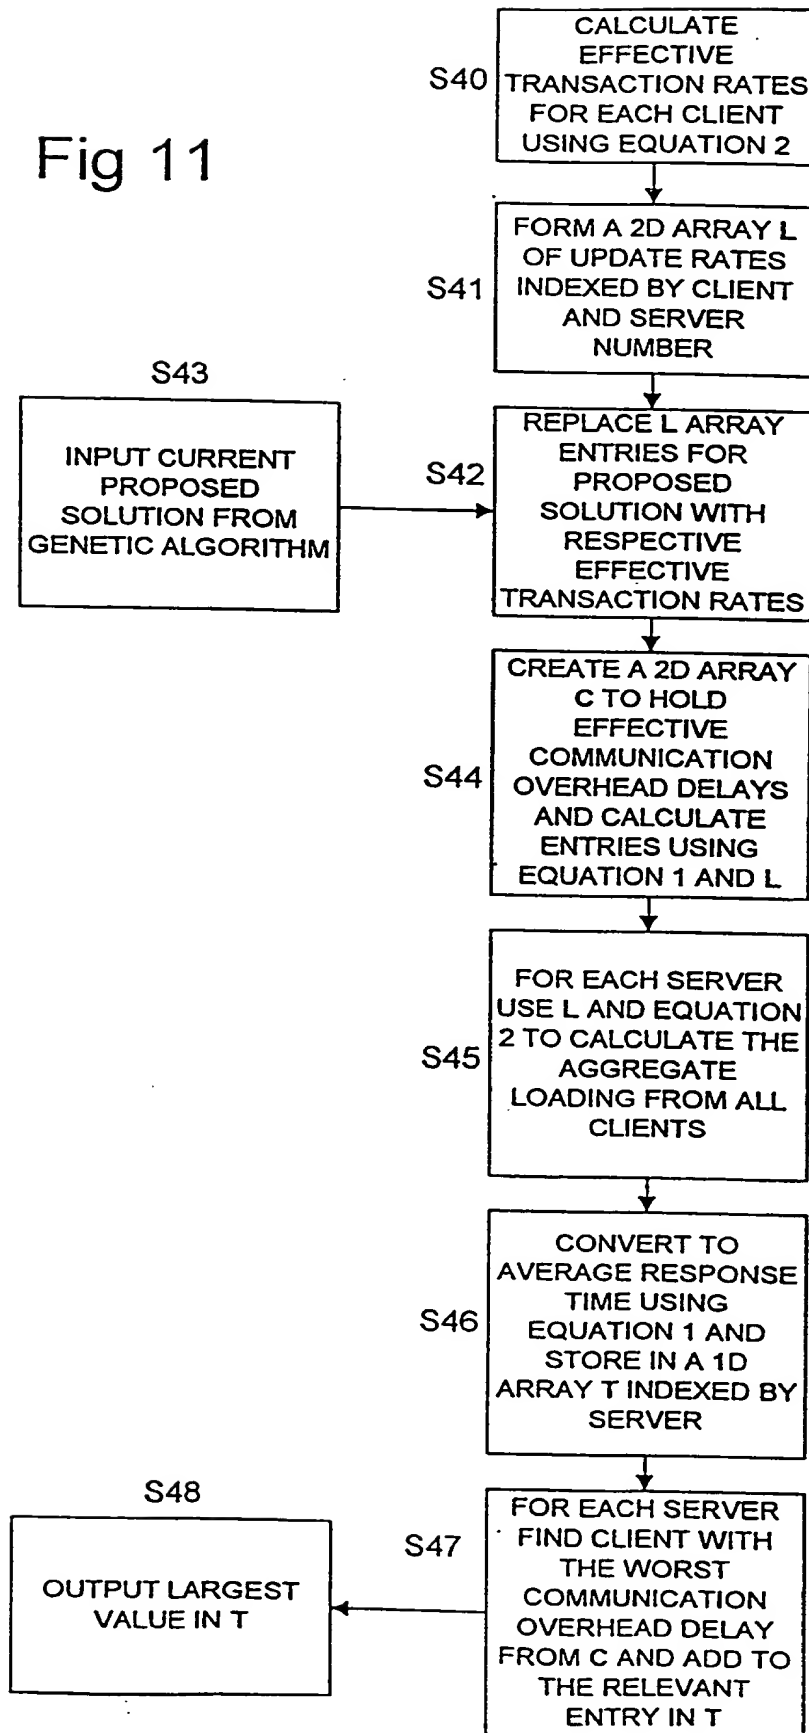


Fig 9

Client	1	2	3	4	5	6	7	8	9	10
Server to be used by client	B	C	A	D	I	G	E	B	A	J

Fig 10

Fig 11



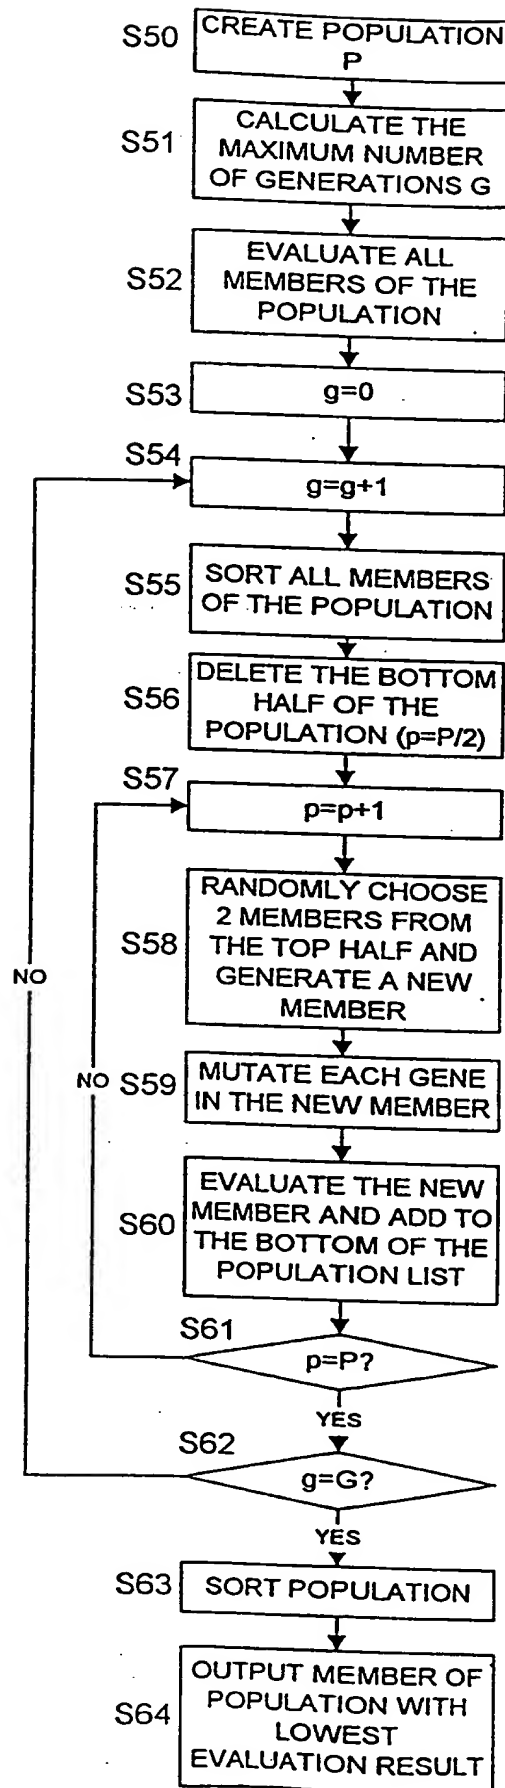
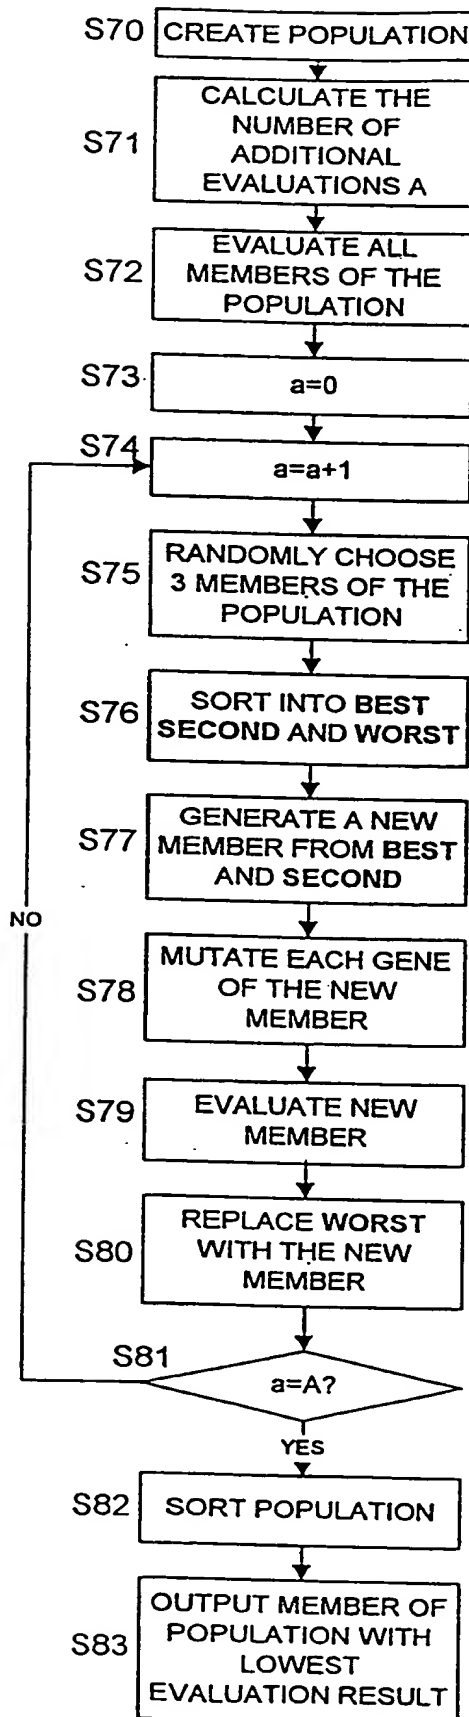


Fig 12

Fig 13



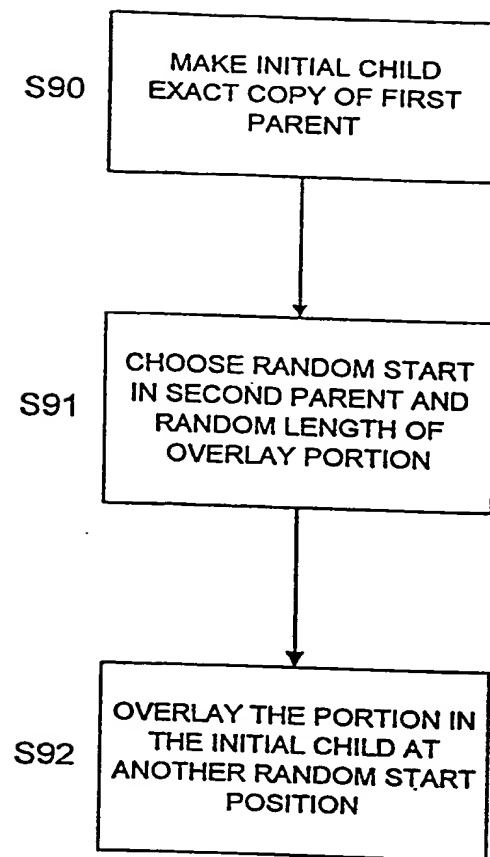


Fig 14

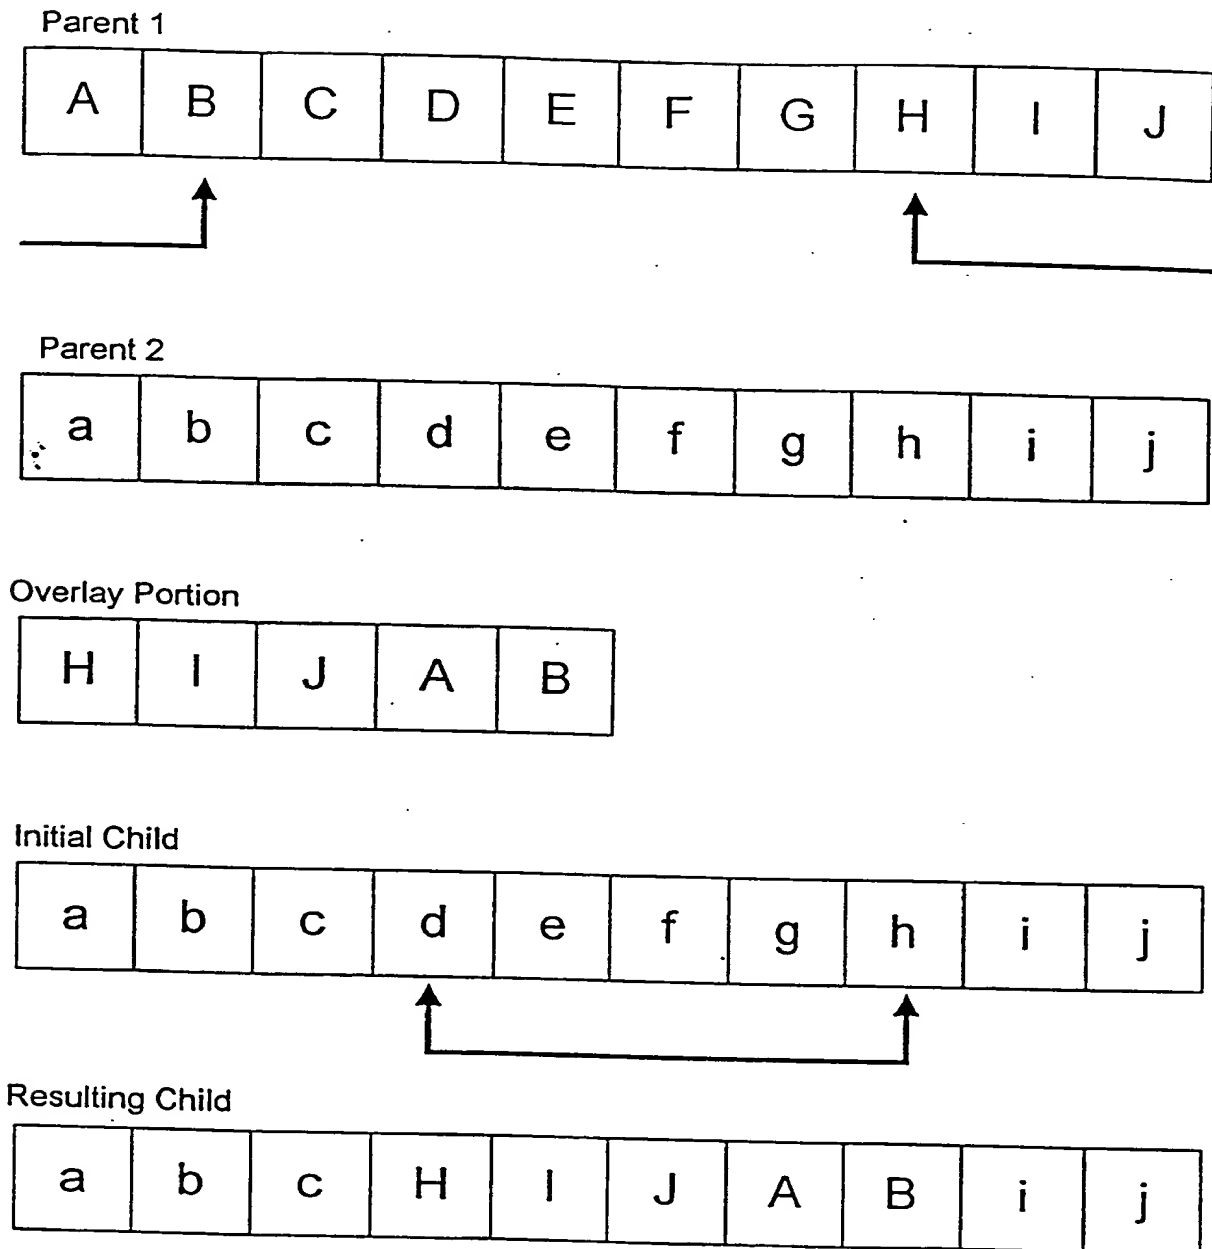


Fig 15




Fig 17

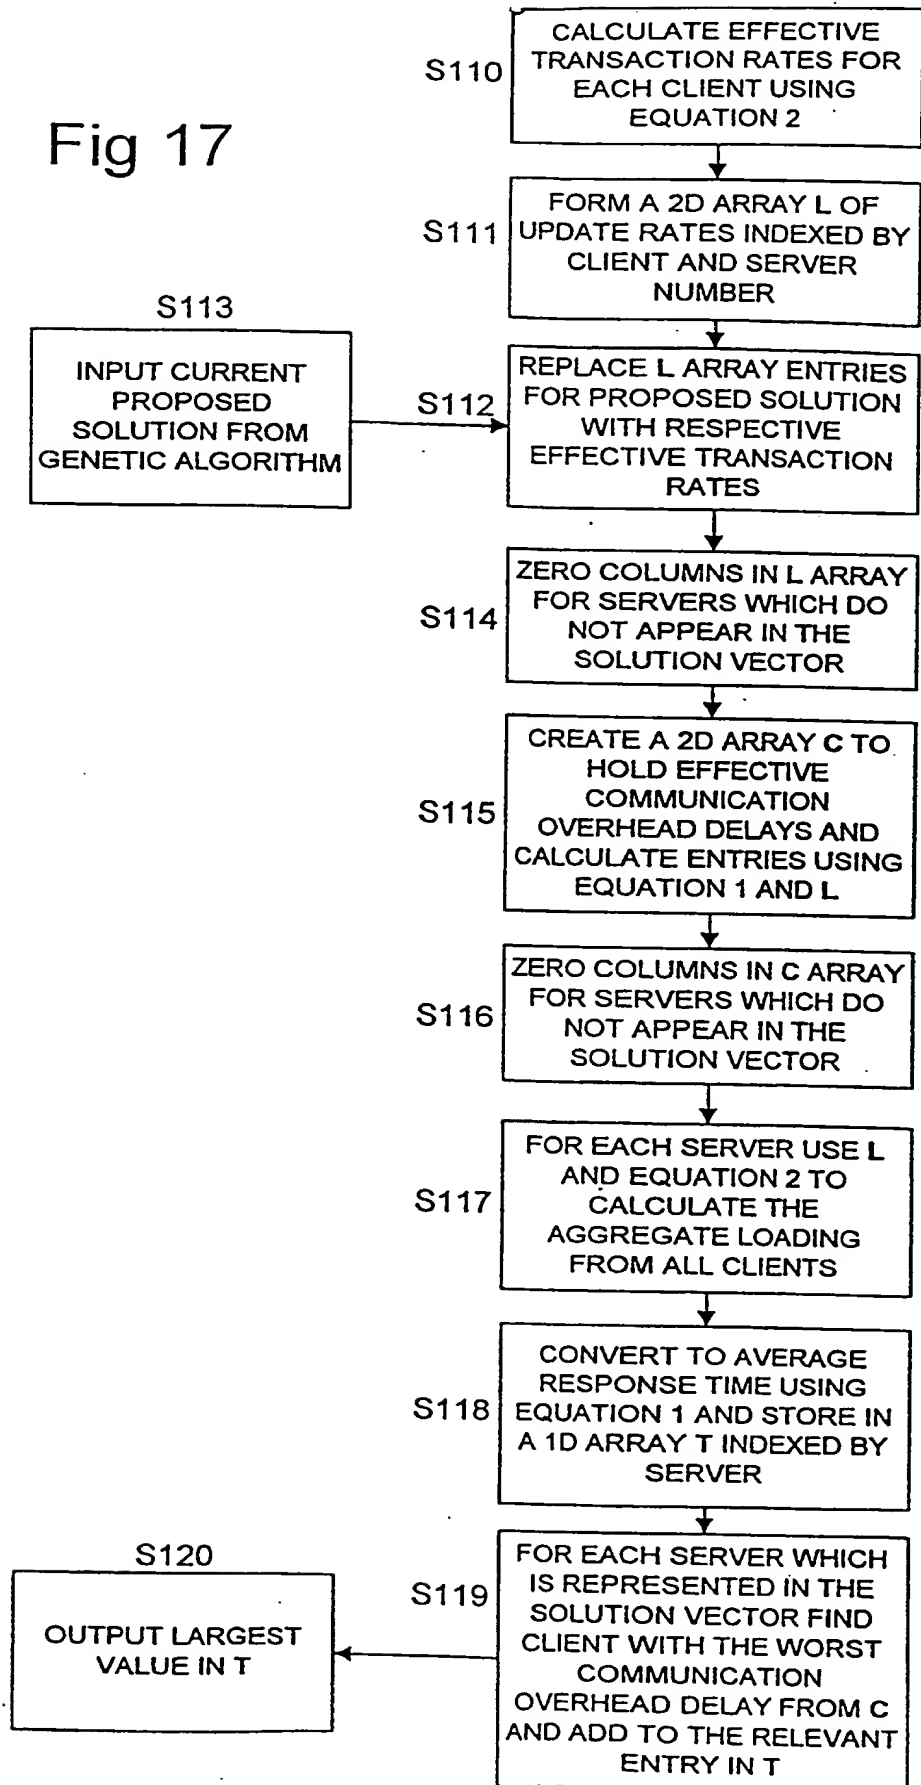


Fig 18

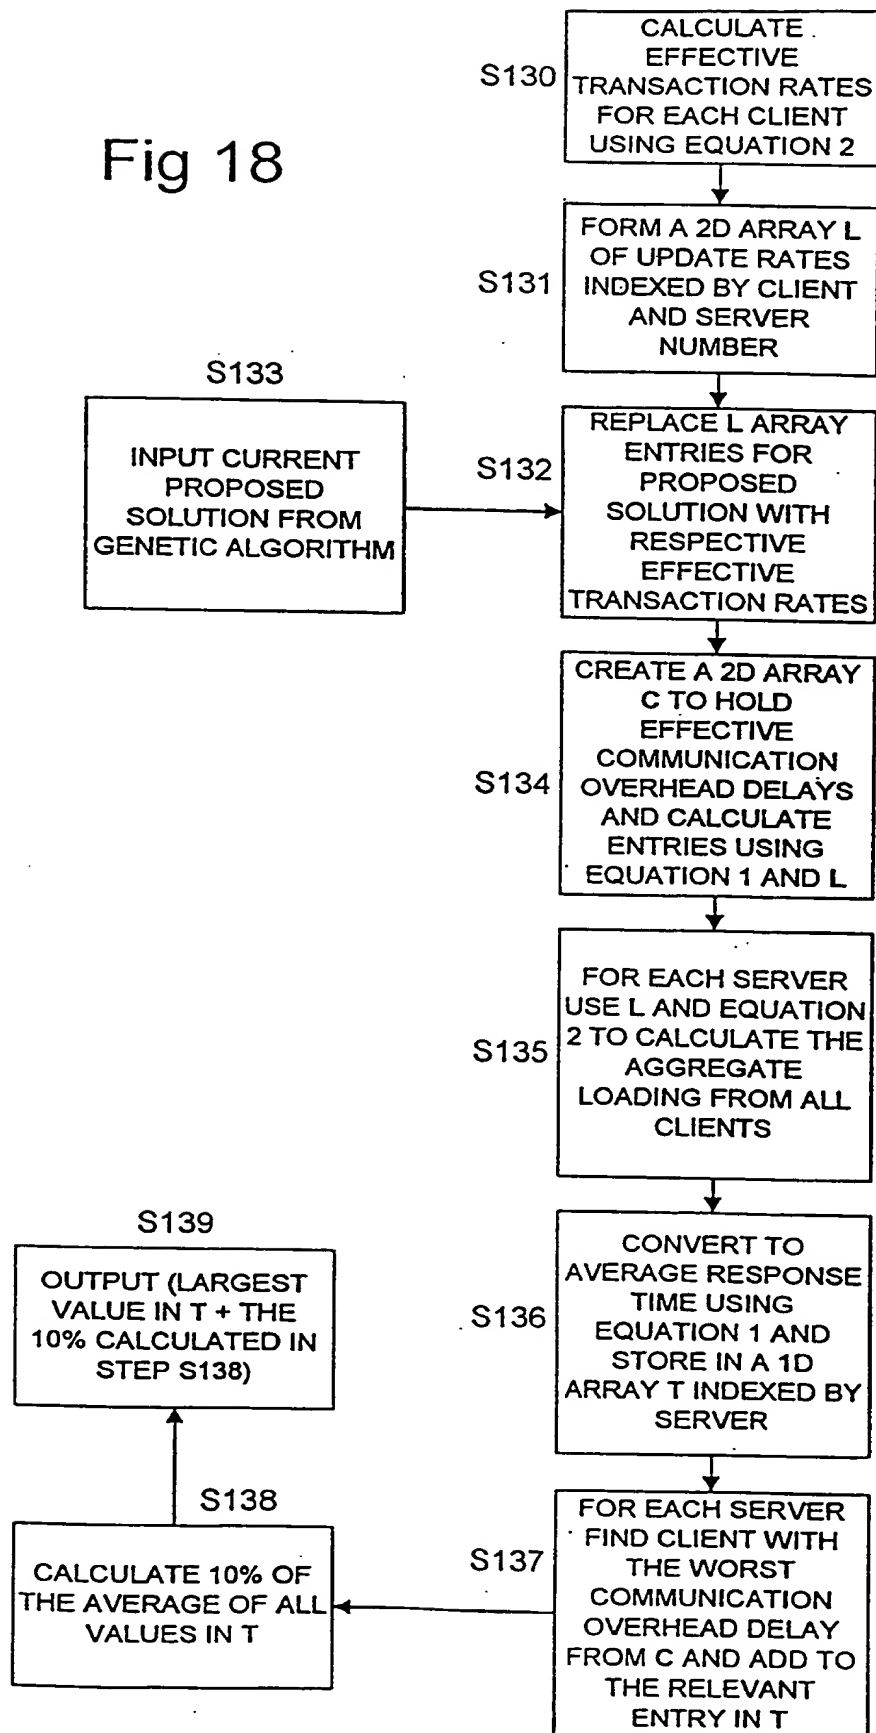


Fig 19

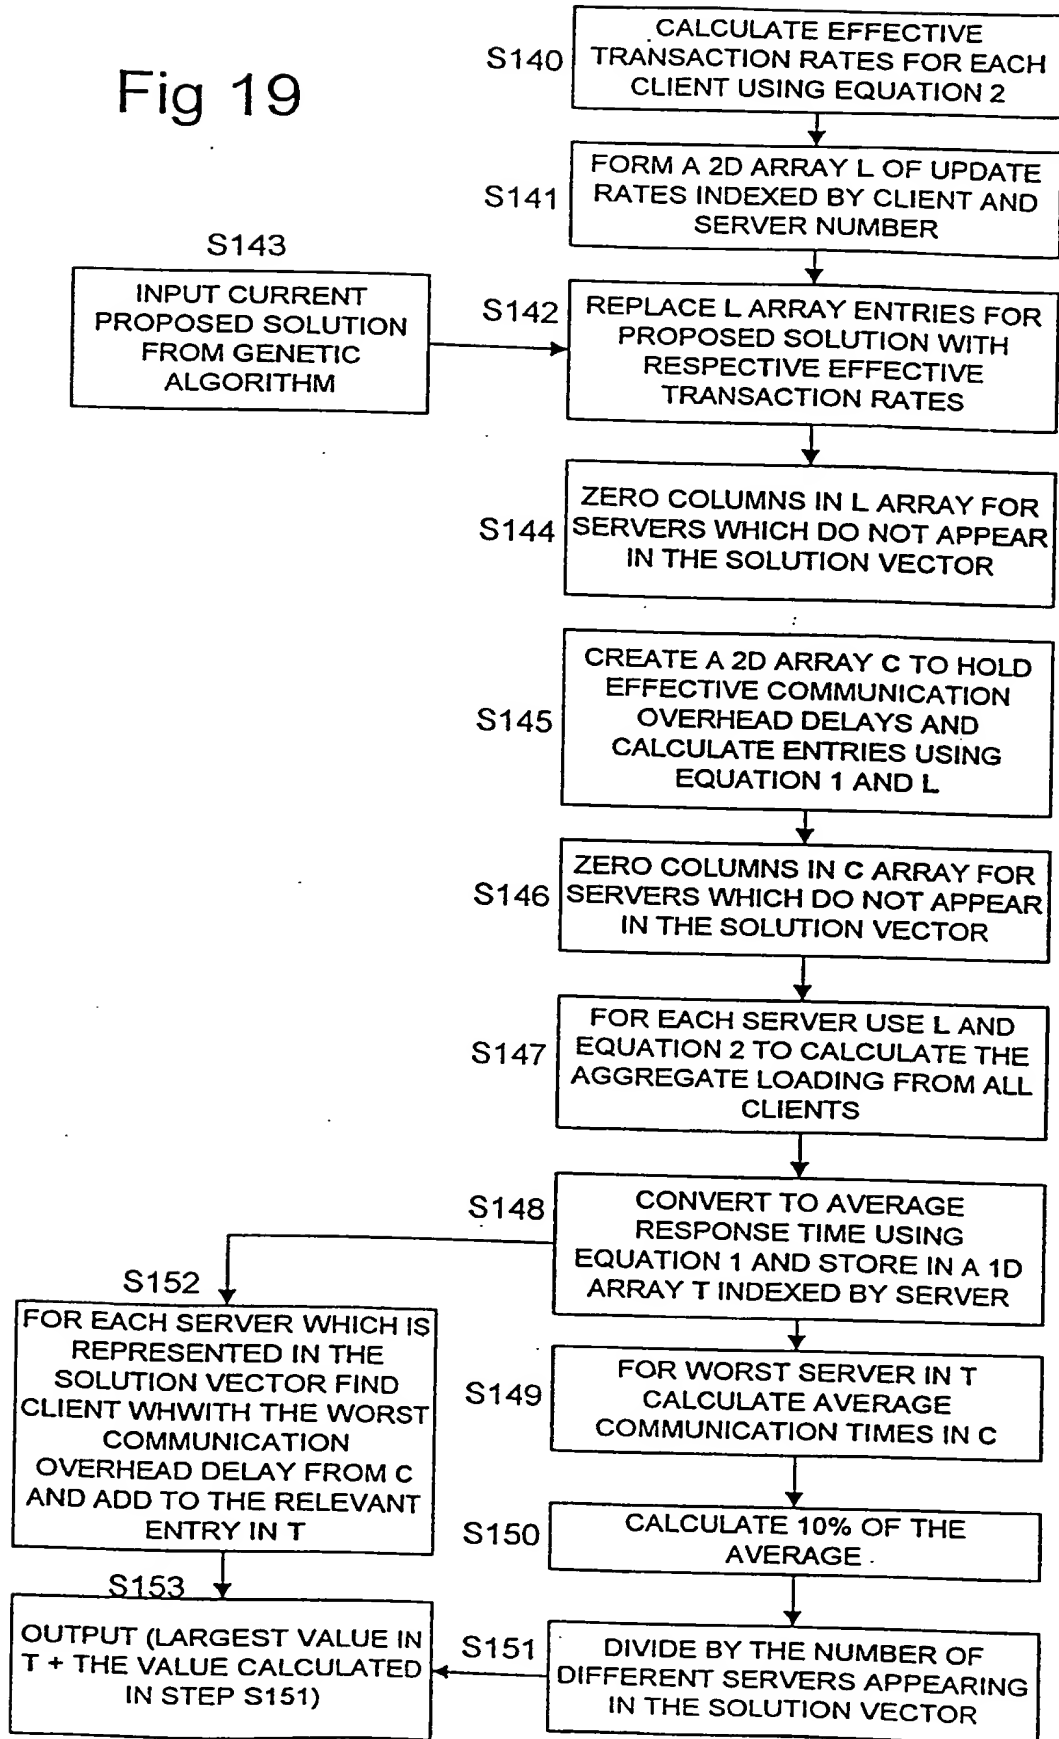


Fig 20

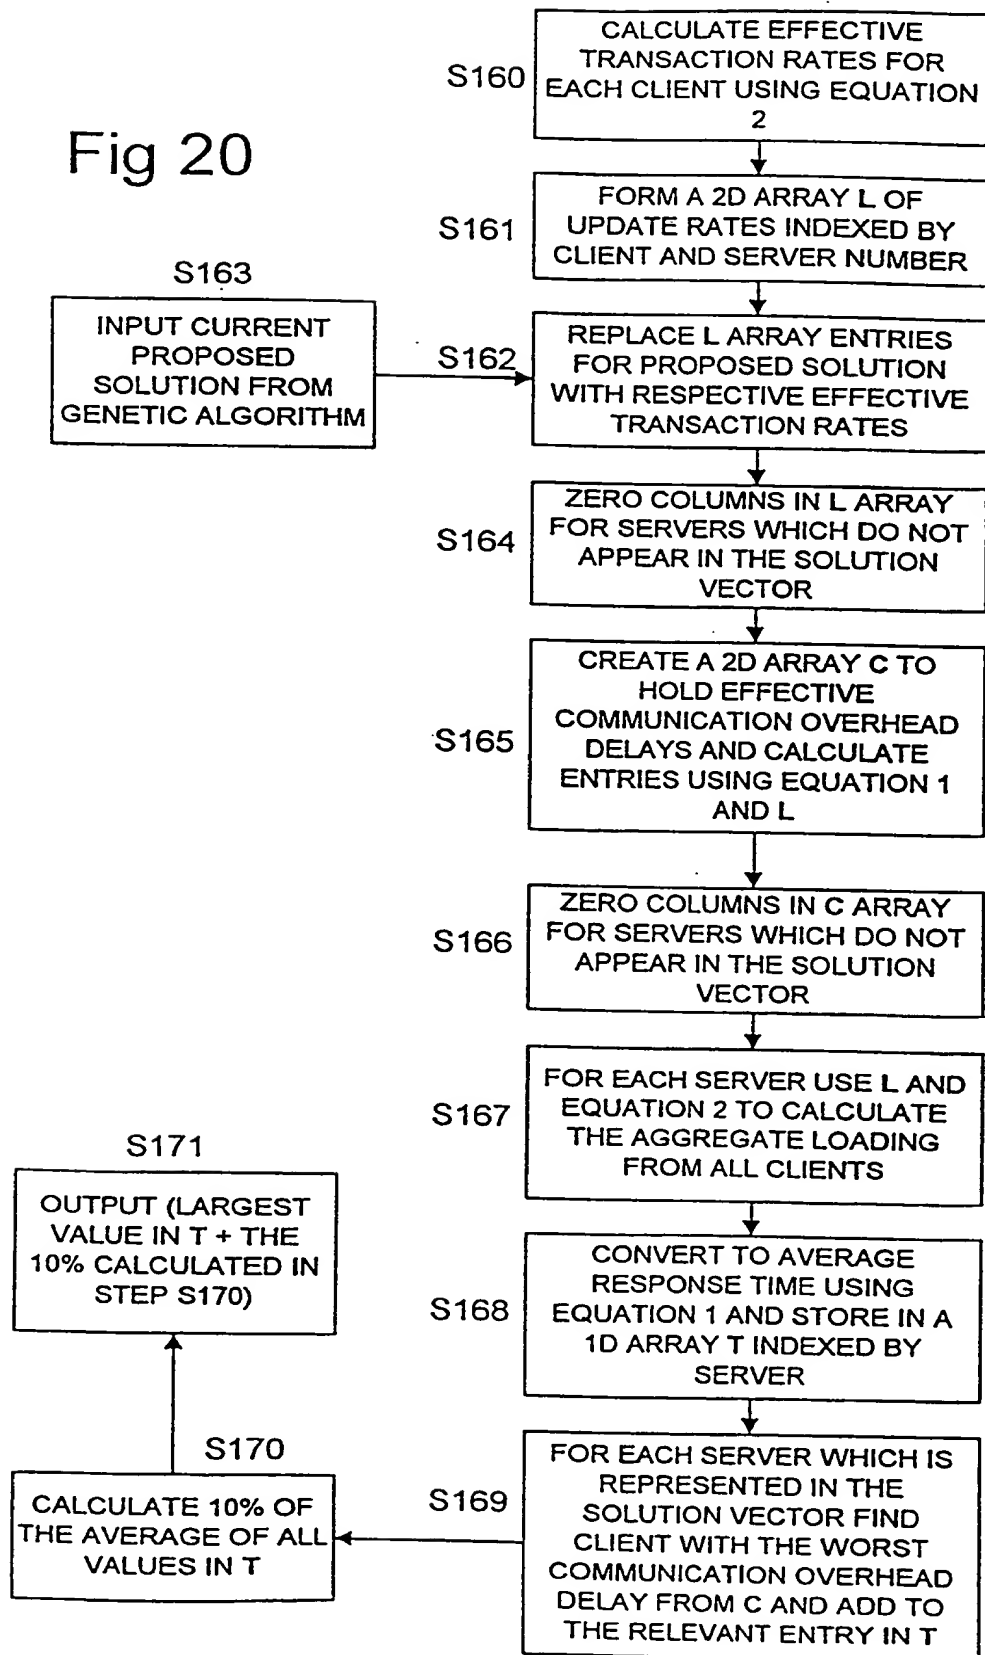


Fig 21

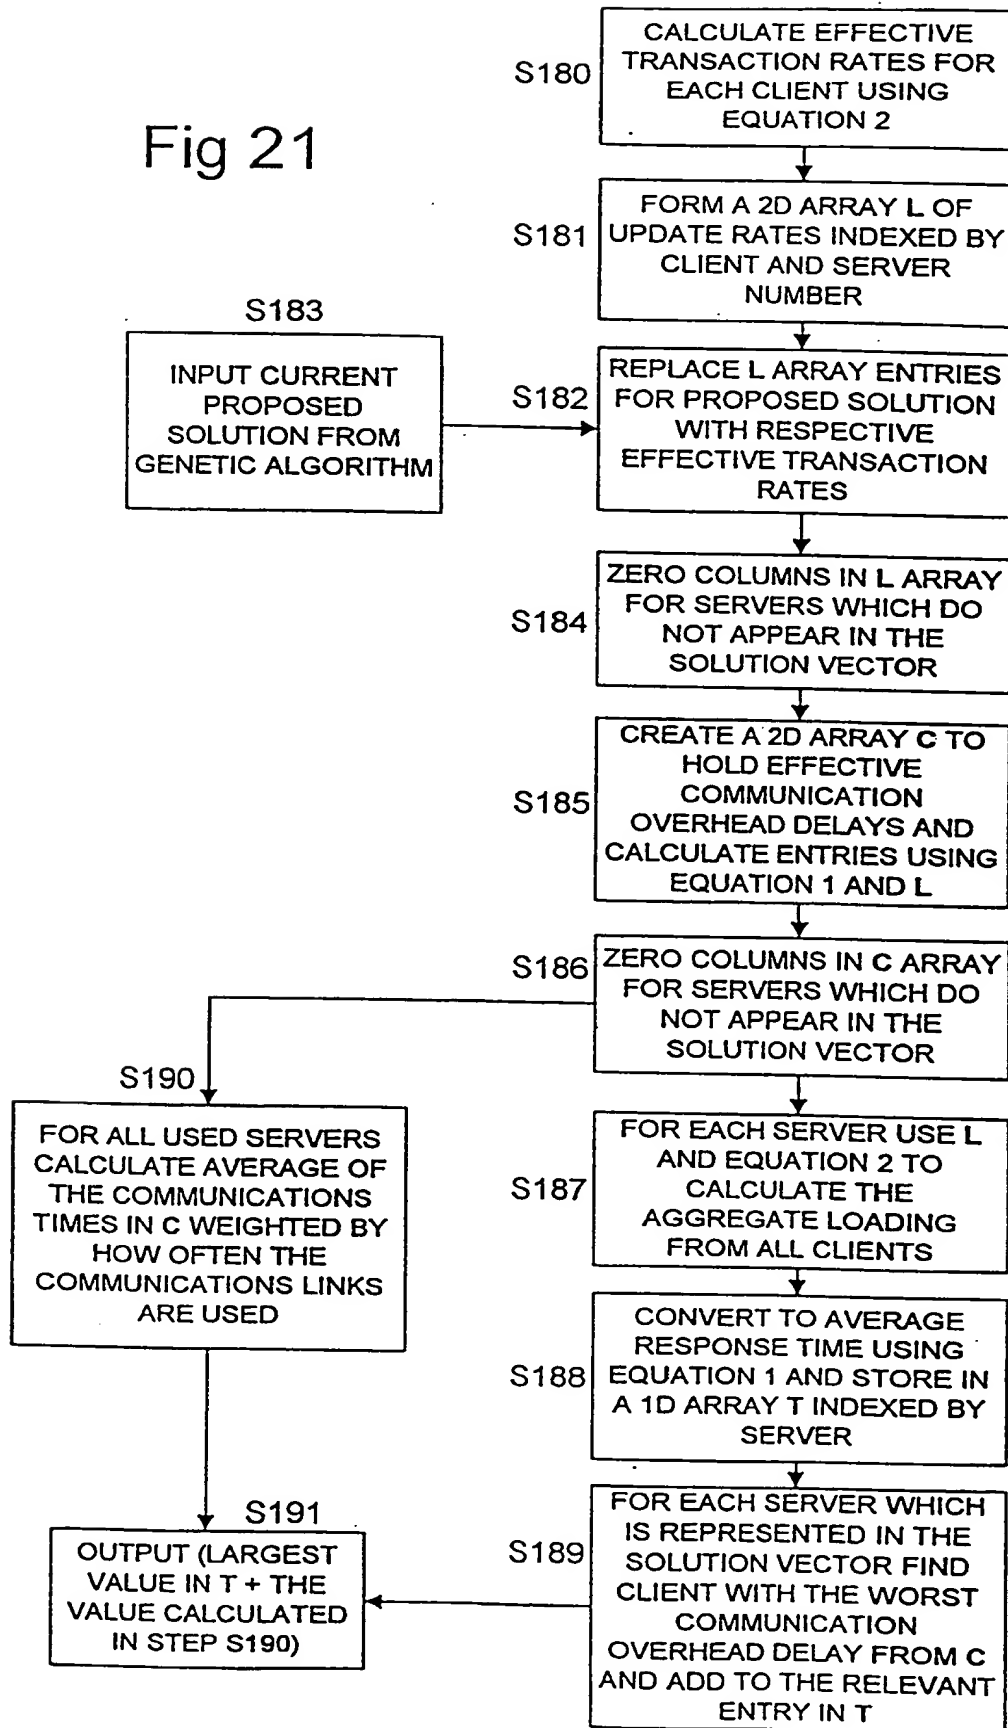


Fig 22

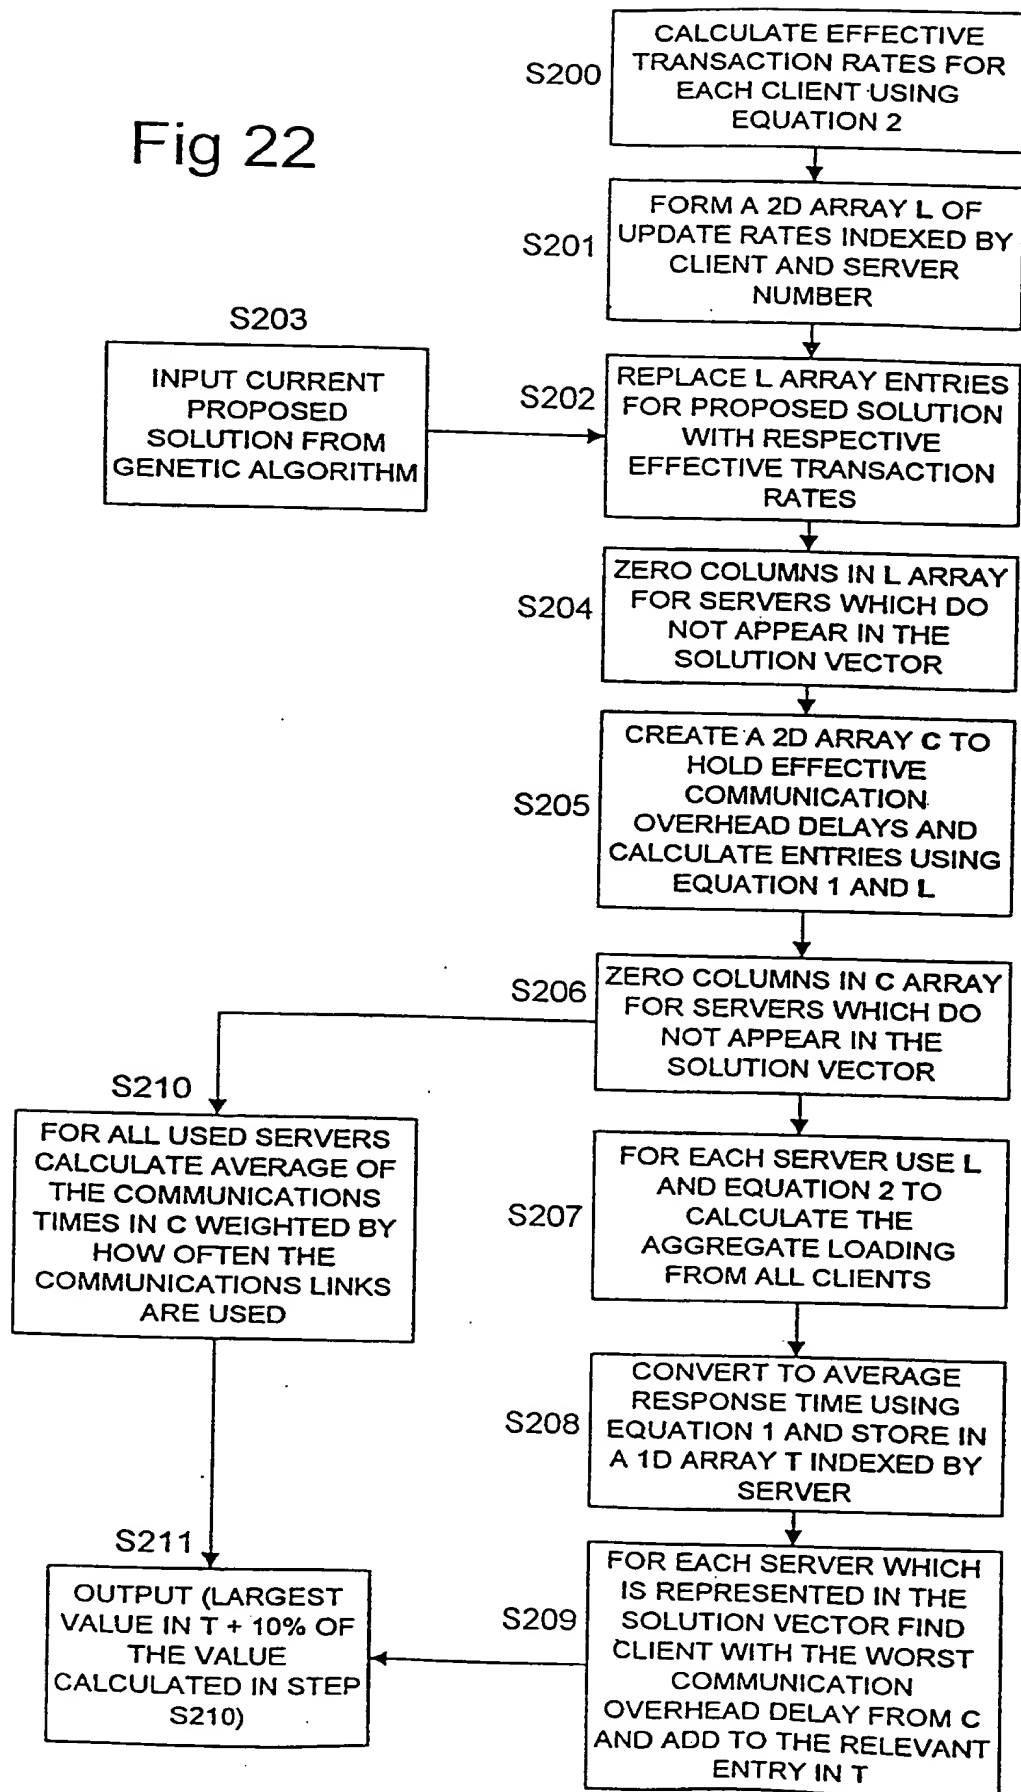


Fig 23

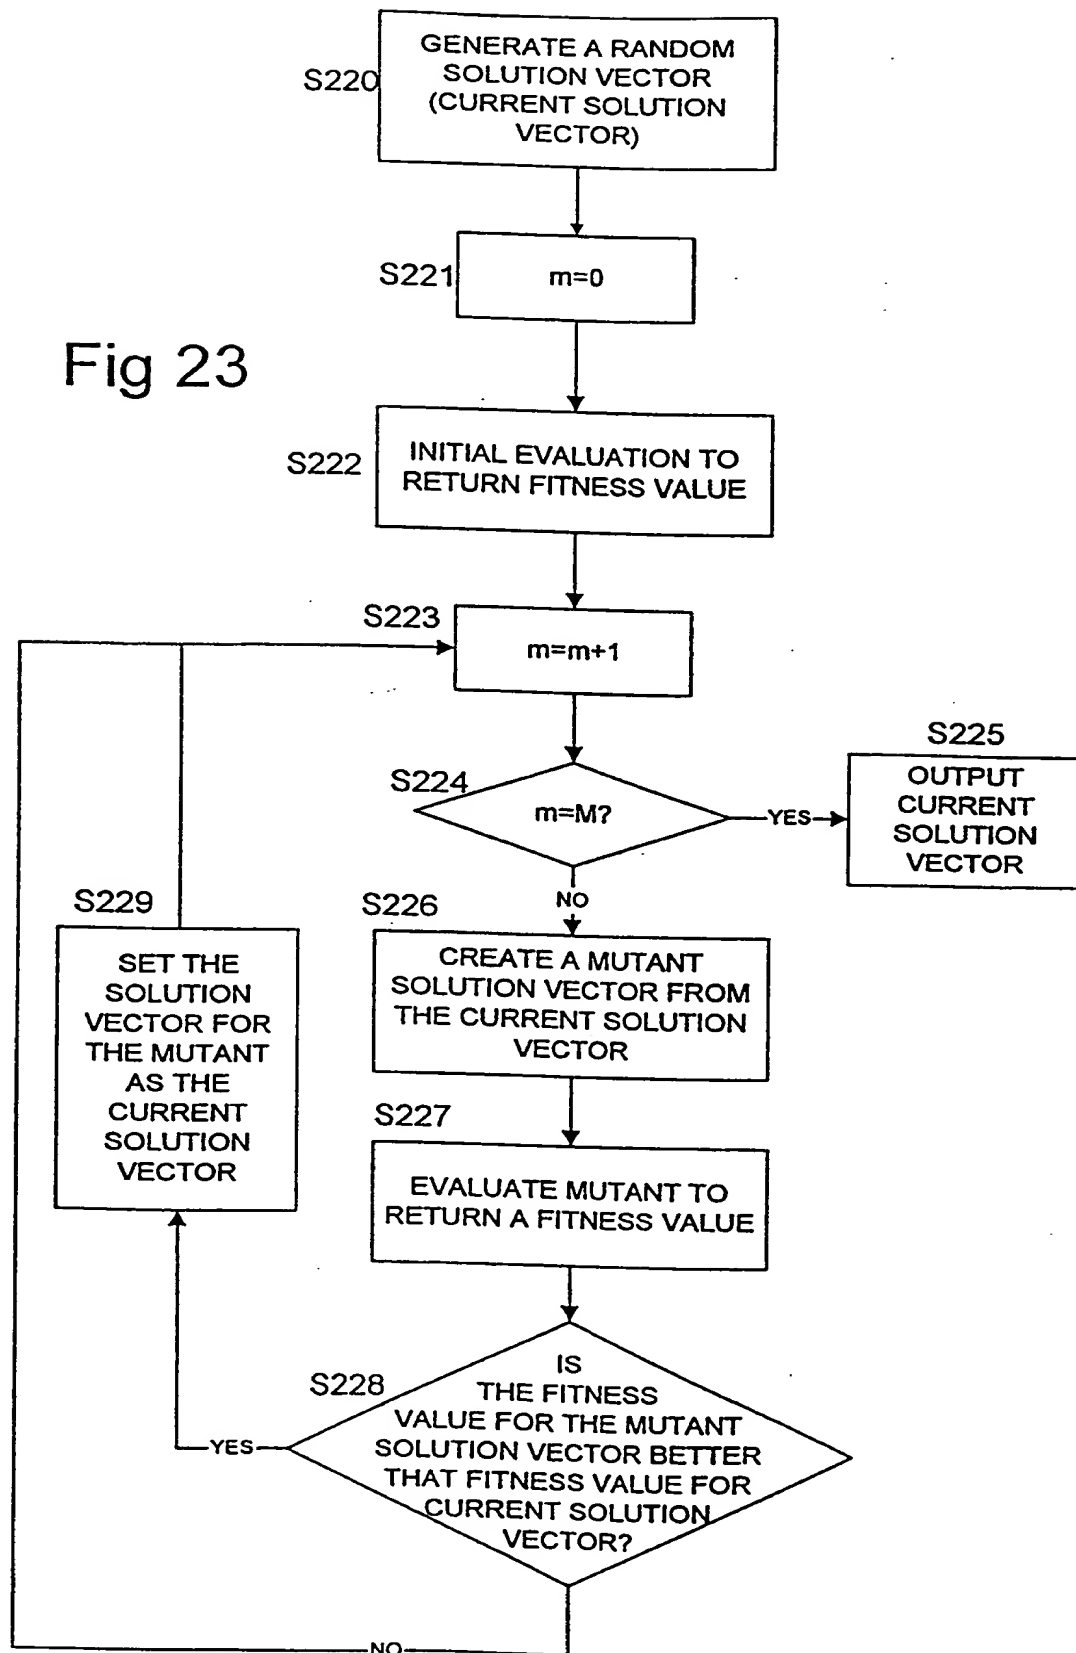
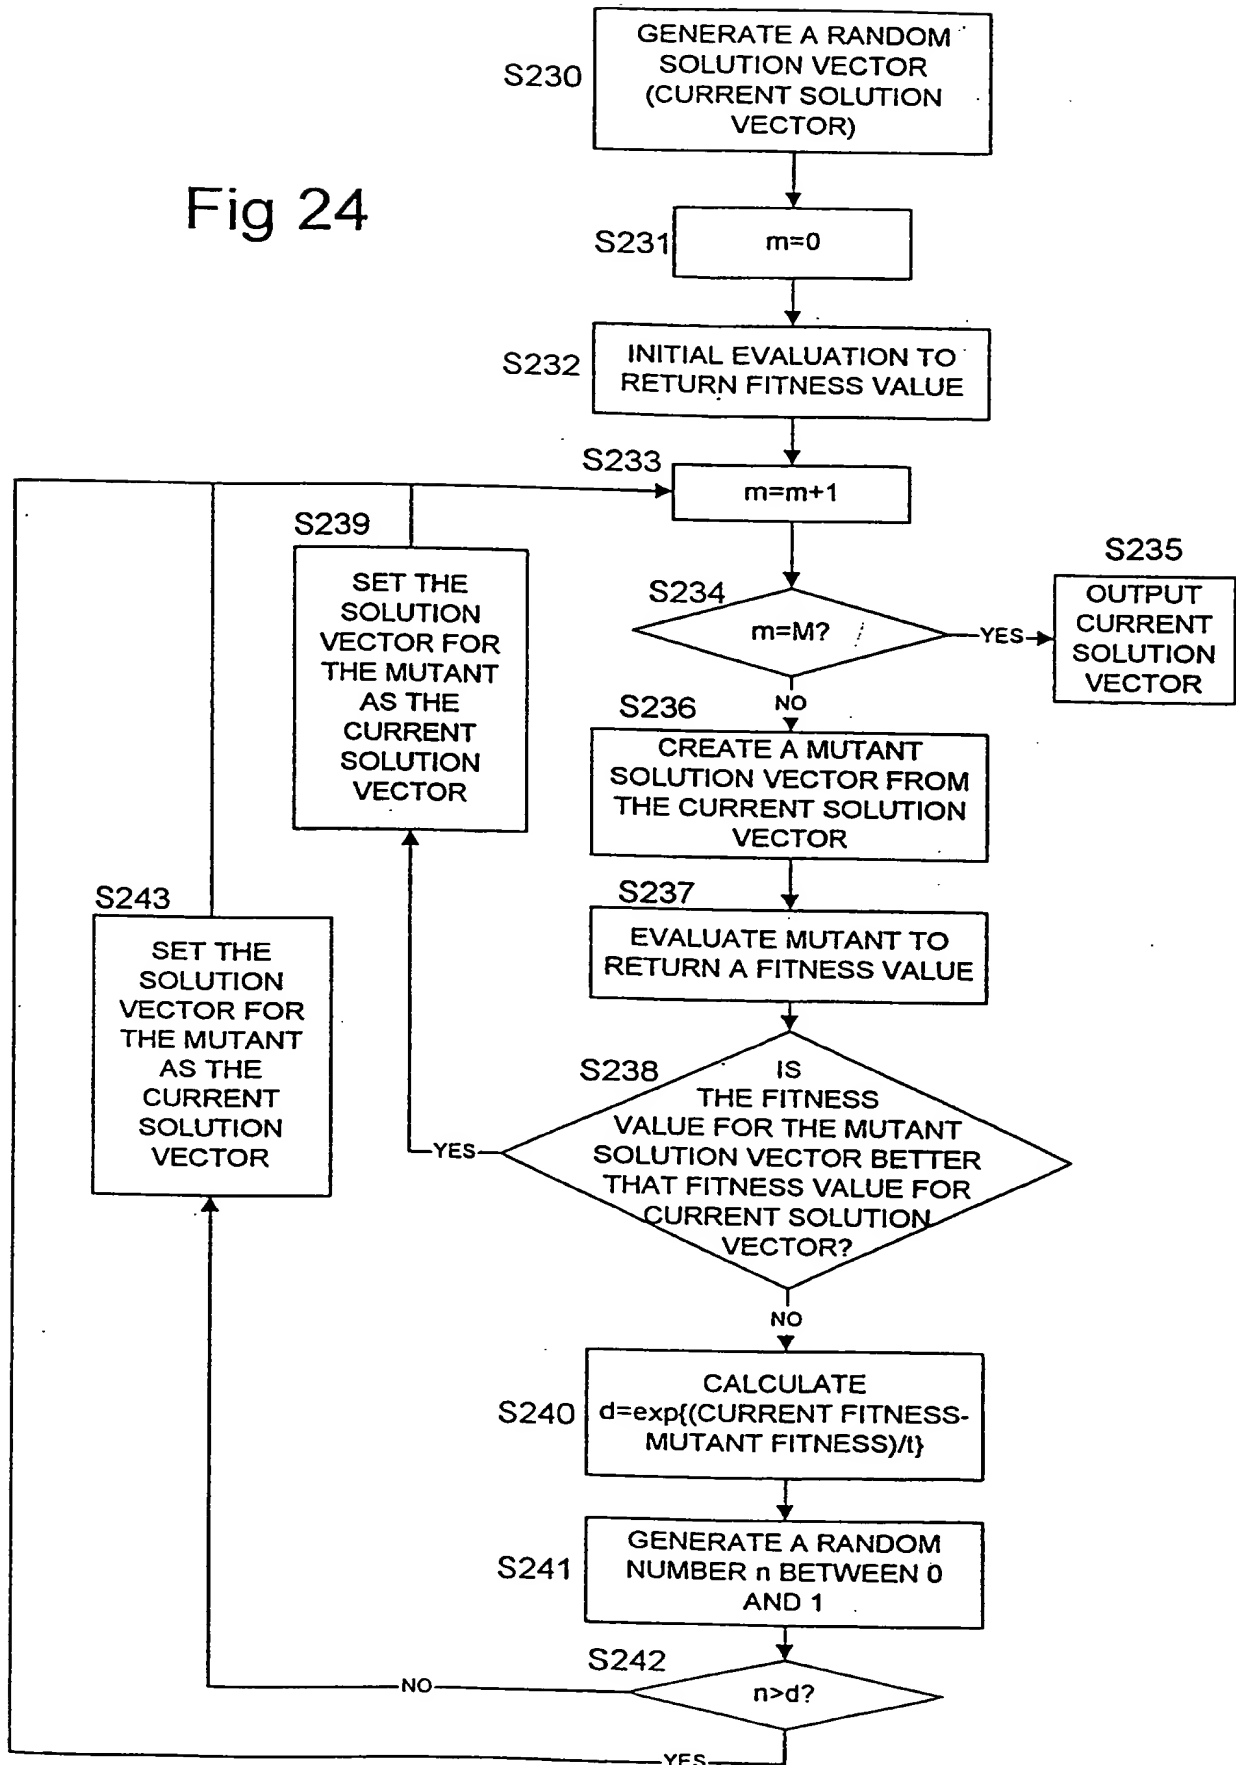


Fig 24



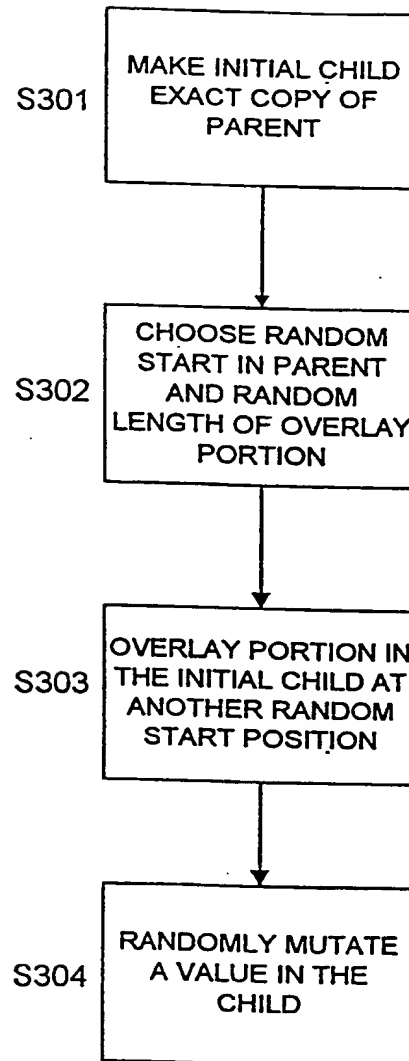


Fig 25

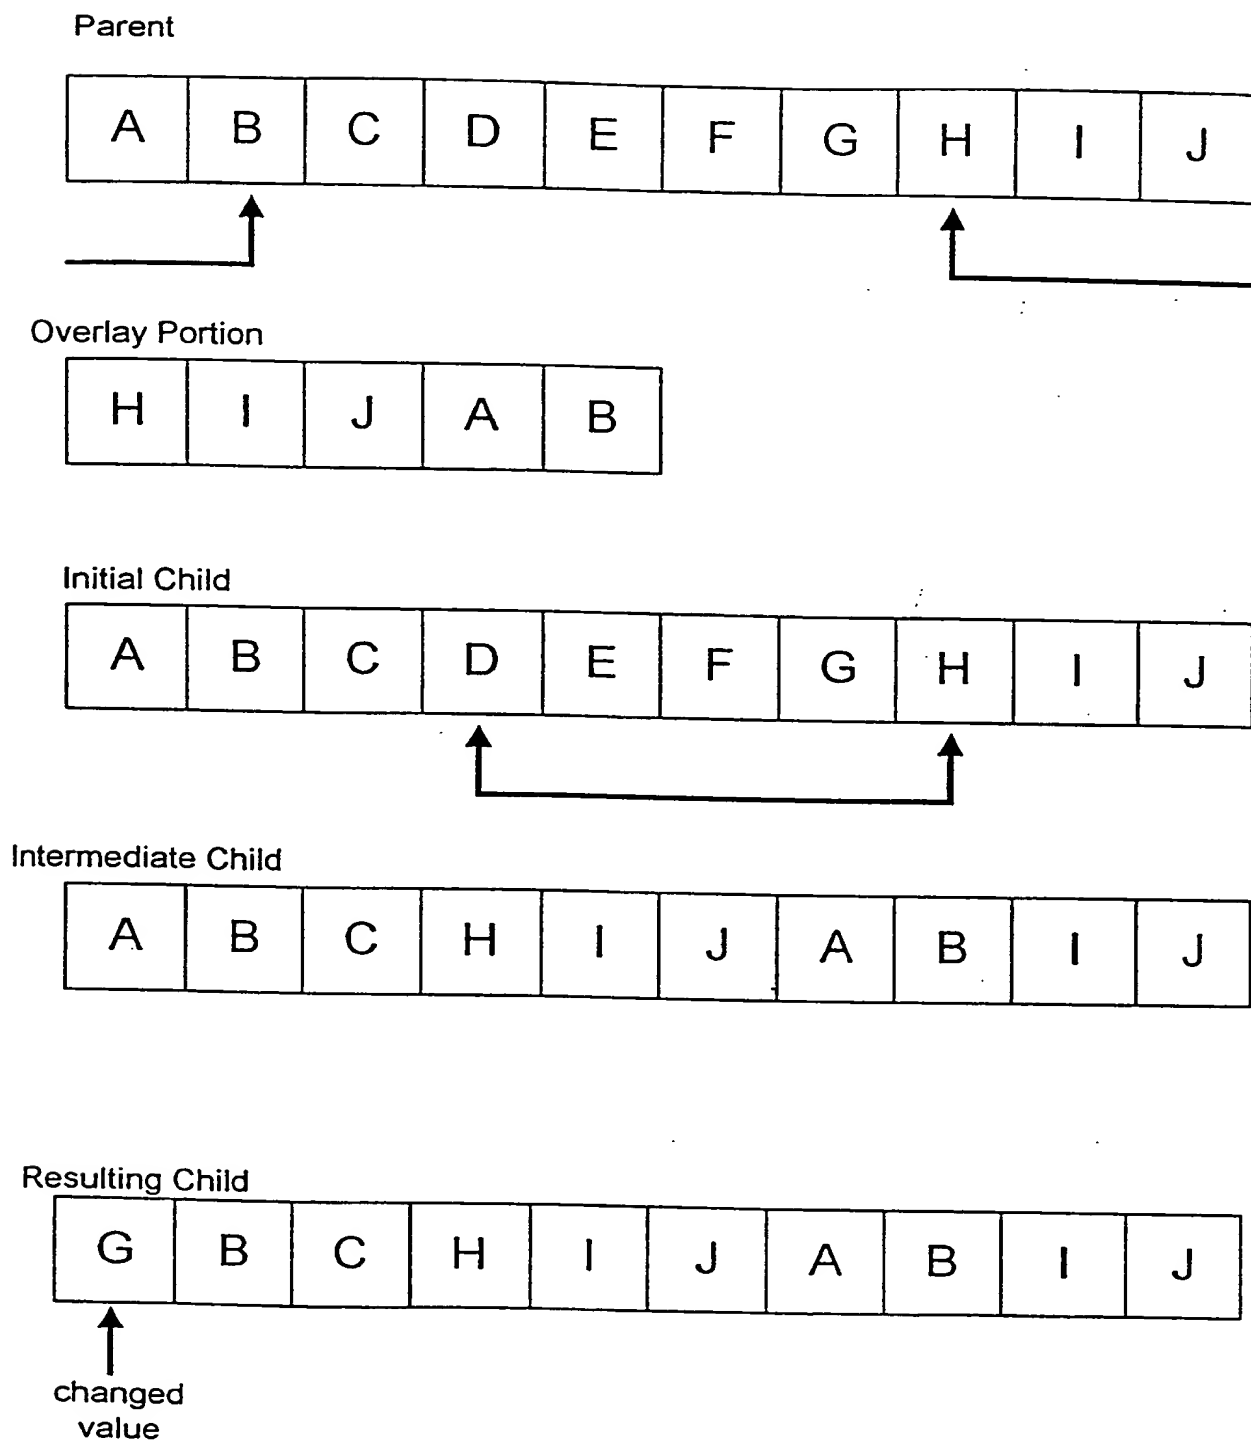


Fig 26




Fig 27




Fig 28




Fig 29




Fig 30

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference A25664 WO	FOR FURTHER ACTION see Notification of Transmittal of International Search Report (Form PCT/ISA/220) as well as, where applicable, item 5 below.	
International application No. PCT/GB 99/ 02449	International filing date (day/month/year) 27/07/1999	(Earliest) Priority Date (day/month/year) 05/08/1998
Applicant BRITISH TELECOMMUNICATIONS PLC et al.		

This International Search Report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This International Search Report consists of a total of 3 sheets.



It is also accompanied by a copy of each prior art document cited in this report.

1. Basis of the report

- a. With regard to the language, the international search was carried out on the basis of the international application in the language in which it was filed, unless otherwise indicated under this item.



the international search was carried out on the basis of a translation of the international application furnished to this Authority (Rule 23.1(b)).

- b. With regard to any nucleotide and/or amino acid sequence disclosed in the international application, the international search was carried out on the basis of the sequence listing:



contained in the international application in written form.



filed together with the international application in computer readable form.



furnished subsequently to this Authority in written form.



furnished subsequently to this Authority in computer readable form.



the statement that the subsequently furnished written sequence listing does not go beyond the disclosure in the international application as filed has been furnished.



the statement that the information recorded in computer readable form is identical to the written sequence listing has been furnished

2. ☐ Certain claims were found unsearchable (See Box I).

3. ☐ Unity of invention is lacking (see Box II).

4. With regard to the title,

the text is approved as submitted by the applicant.



the text has been established by this Authority to read as follows:

5. With regard to the abstract,

the text is approved as submitted by the applicant.



the text has been established, according to Rule 38.2(b), by this Authority as it appears in Box III. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority.

6. The figure of the drawings to be published with the abstract is Figure No.

as suggested by the applicant.



because the applicant failed to suggest a figure.



because this figure better characterizes the invention.

3

None of the figures.

INTERNATIONAL SEARCH REPORT

Intern Application No

PCT/GB 99/02449

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F17/30 G06F15/18 G06F11/34 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>GOLDBERG D E: "Genetic Algorithms in Search, Optimization, and Machine Learning" January 1989 (1989-01) , ADDISON WESLEY , USA XP000667450</p> <p>page 6 -page 10 page 61 -page 82 page 100 page 115 page 119 page 208 -page 212</p> <p>--- -/--</p>	<p>1-4,7, 10-13, 16, 48-51, 54, 57-60, 63,95,96</p>

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

22 October 1999

Date of mailing of the international search report

05/11/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Dupuis, H

INTERNATIONAL SEARCH REPORT

Intern al Application No
PCT/GB 99/02449

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>PARK SEONG-JIN ET AL: "Data allocation considering data availability in distributed database systems"</p> <p>PROCEEDINGS OF THE 1997 INTERNATIONAL CONFERENCE ON PARALLEL AND DISTRIBUTED SYSTEMS; SEOUL, SOUTH KOREA DEC 10-13 1997, 1997, pages 708-713, XP002093736</p> <p>Proc Int Conf Parallel Distrib Syst</p> <p>ICPADS; Proceedings of the Internatoinal Conference on Parallel and Distributed Systems - ICPADS 1997 IEEE Comp Soc, Los Alamitos, CA, USA</p> <p>abstract</p> <p>paragraph '0001!</p> <p>paragraph '0003!</p> <p>---</p>	<p>20, 35, 67, 82, 95, 96</p>
A	<p>GREGORY MICHAEL: "Genetic algorithm optimization of distributed database queries"</p> <p>PROCEEDINGS OF THE 1998 IEEE INTERNATIONAL CONFERENCE ON EVOLUTIONARY COMPUTATION, ICEC'98; ANCHORAGE, AK, USA MAY 4-9 1998, 1998, pages 271-276, XP002093737</p> <p>Proc IEEE Conf Evol Comput Proc</p> <p>ICEC; Proceedings of the IEEE Conference on Evolutionary Computation, ICEC 1998 IEEE, Piscataway, NJ, USA</p> <p>---</p>	<p>1, 10, 20, 35, 48, 57, 67, 82</p>
A	<p>SINCLAIR M C: "The application of a genetic algorithm to trunk network routing table optimisation"</p> <p>TENTH UK TELETRAFFIC SYMPOSIUM.</p> <p>PERFORMANCE ENGINEERING IN TELECOMMUNICATIONS NETWORK, MARTLESHAM HEATH, UK, 14-16 APRIL 1993, pages 2/1-6, XP002093738</p> <p>1993, London, UK, IEE, UK</p> <p>paragraph '0001!</p> <p>paragraph '05.1!</p> <p>-----</p>	<p>1-5, 10-14, 48-52, 57-61</p>